What Do Measures of Decision-Making Ability Tell Us About Preventive Health Behavior? Evidence from Flu and COVID-19 Vaccination

By Xuqian Ma*

November 4, 2026

Vaccination is a choice under uncertainty that draws on two different decision making abilities: forming beliefs about risks and making fewer mistakes when making the same kind of choice again and again. Using the Understanding America Study, I measure the first with cognitive tests of quantitative ability and the second with choice consistency in an incentivized revealed-preference experiment. For influenza (2014–2026), neither predicts vaccination on average, but receipt falls over time especially post COVID among the least consistent adults, most sharply among Republicans, and the consistency gap grows most among adults with fewer chronic conditions, groups likely to perceive less risk. For COVID-19 (2021), quantitative ability predicts who got vaccinated earlier. Higher-ability adults held more favorable beliefs about vaccine safety and effectiveness before most could get a dose, and these beliefs account for about 60 percent of the gap. After adjusting for baseline beliefs, lower-ability adults more often reported not knowing how to book or where to go. This points to prosocial and norm-based messages from messengers people trust for familiar vaccines, and to early, clear safety information and appointment guidance for new ones.

Keywords: quantitative ability; choice consistency; revealed preference; cognitive tests; vaccination

JEL codes: D81; D91; I12; I14; J14

*University of California, Berkeley, Department of Health Policy and Management. ma000092@berkeley.edu. I am grateful to William H. Dow, Shachar Kariv, Edward Miguel, Rebecca Staiger, and Dan Zeltzer for their helpful comments. The project uses data from surveys administered by the Understanding America Study, which is maintained by the Center for Economic and Social Research (CESR) at the University of Southern California. The content of this paper is solely the responsibility of the author and does not necessarily represent the official views of USC or UAS.

1 Introduction

Vaccination is a private decision with public consequences. In economic epidemiology, a person weighs her perceived risk of illness and the protection a dose offers against its side effects and the cost of getting it (Philipson, 2000). Disease acts like a tax on exposure, inducing costly prevention (Philipson, 1995), so the demand for vaccination rises with the prevalence of disease and falls as the disease recedes (Philipson, 1996). Each dose also protects others (Boulier et al., 2007), so private demand falls short of what society would choose, and the private gain shrinks as prevalence falls (Geoffard and Philipson, 1997). Voluntary coverage is therefore set by the adults whose expected gain from a dose barely exceeds its cost, and small changes in how they judge that gain, or in how often they make mistakes in deciding, move coverage.

Since 2020, the decision has become harder in two ways. For a new vaccine, adults had to judge the risks in a polarized information environment: misinformation about the virus spread widely on social media (Cinelli et al., 2020), exposure to false claims about the vaccines lowered stated intentions to take them (Loomba et al., 2021), and responses to public-health guidance divided along party lines (Allcott et al., 2020). For a familiar vaccine, the private gain came to look smaller to some adults. Attitudes toward the COVID-19 vaccines spilled over to influenza, turning more negative on the political right (Lunz Trujillo et al., 2024; LaCour and Bell, 2024), and reported receipt or plans for a flu shot declined between 2020 and 2024 (Melchinger et al., 2025). Among Republicans, the predicted probability of a flu shot fell from 0.54 before the pandemic to 0.44 after the COVID-19 vaccine rollout, against 0.63 and 0.61 among Democrats (Sandlin, 2025). Federal advice has also moved from routine COVID-19 vaccination (Panagiotakopoulos et al., 2024) to individual-based decision-making (Centers for Disease Control and Prevention, 2025), which leaves more of this weighing to each adult.

The two difficulties call on different decision making abilities. Forming beliefs about protection and side effects matters most for a new vaccine whose evidence is disputed: within each party, adults with greater cognitive sophistication held fewer misperceptions about COVID-19 and were less hesitant to vaccinate (Pennycook et al., 2022). The second ability is making fewer mistakes when making the same kind of decision again and again. Many health behaviors, an annual flu shot among them, require people to keep making good choices, and the cost of each mistake adds up. Mistakes are also most likely when the expected gain is small, so when a familiar vaccine comes to look less worthwhile, an adult who values the shot but pays less attention to her choices is more likely to skip it in some seasons. Coverage thus depends on how individual adults make these decisions, and on whether health systems

can tell which of them are falling behind.

Most studies of how adults make this decision use cognitive tests, which mainly capture the first decision making ability. Test scores predict vaccine hesitancy and the speed of COVID-19 vaccination (Batty et al., 2021; Halstead et al., 2022; Elinder et al., 2023). In Swedish registers, the highest-scoring conscripts reached 80 percent COVID-19 vaccination coverage about four months before the lowest-scoring, and pre-booked appointments narrowed the gap (Elinder et al., 2023). A test score, however, records answers graded against a key. It tells us how well a person handles numerical information, but not how often she makes mistakes when she faces the same kind of choice many times.

Revealed-preference theory offers a measure of the second decision making ability that does not need a correct answer. In an incentivized experiment, a person makes many choices of the same kind, and the measure records how far those choices fall short of being explained by any stable preference (Afriat, 1973; Varian, 1982; Choi et al., 2014). A shortfall arises when one choice contradicts another, which is a mistake by the person's own standard. Two people who would never agree about how much risk to take can therefore both attain the highest score. I call this measure choice consistency. We do not know whether it predicts who gets vaccinated and who does not.

This paper compares the two measures for the same U.S. adults in the Understanding America Study (UAS), a probability-based internet panel (Alattar et al., 2018). Between 2014 and 2019, spanning UAS–HRS Waves 1 and 2, respondents completed three cognitive assessments and a financial-literacy test, which I combine into a measure of quantitative ability, and an incentivized choice experiment that measures choice consistency. I link the two measures to self-reported flu vaccination in five UAS–HRS waves from 2014 to 2026 and to COVID-19 vaccination in ten survey waves during the 2021 rollout. Following the same people through a familiar decision and a new one shows whether each measure predicts a different vaccination path.

For influenza, neither measure predicts vaccination on average, but the choice-consistency association rises over time and is larger after 2020 than before, a change I call the flu divergence. The rise comes from falling receipt among the least consistent. Between 2014–2017 and 2025–2026, adjusted vaccination falls by 9.4 percentage points in the bottom consistency tertile and changes little in the top tertile. In both the bottom and top quantitative-ability tertiles, it falls by about 4 points. Some of this separation already appears in 2017–2019. Between 2014–2017 and the pooled 2021–2026 periods, receipt in the bottom consistency tertile falls by 15.3 percentage points among Republicans and changes little among Democrats. The gap between the top and bottom tertiles grows by 11.4 points among respondents with fewer than two chronic conditions and by 1.2 points among those with more. These are groups

for whom the expected gain from a flu shot is likely small or fell after 2020. So the patterns are suggestive evidence that adults who make more mistakes across repeated decisions, and who gain little from a flu shot, become less likely to get vaccinated over time. This points to messages about protecting others and social norms, which do not rely on a person's own perceived risk, delivered by messengers these adults trust.

During the COVID-19 rollout, neither measure predicts receipt at first. Quantitative ability then becomes more positively associated with receipt as eligibility expands, reaching about 4.7 percentage points per standard deviation during broad adult availability, roughly a quarter of the college gap in that phase. Higher-ability respondents already held more favorable beliefs about vaccine safety and effectiveness when the rollout began, and these beliefs changed little over time. Adding them to the model reduces the later widening in receipt by about 60 percent; adding institutional trust reduces it by about a quarter, and adding perceived infection risk does not reduce it. After adjusting for Wave 21 beliefs, lower-ability respondents more often reported not knowing how to book or where to go (by 2.8 percentage points; $p = 0.064$), and most who gave only this reason were vaccinated within a few months. These patterns are suggestive evidence that higher-ability adults formed more favorable beliefs about the new vaccines before they could get a dose and acted on them once access opened, with a smaller and less precise role for knowing how to get a dose. They motivate testing clear, early safety communication, along with help getting a dose such as pre-booked appointments (Elinder et al., 2023).

The paper makes two contributions. First, to my knowledge, it is the first to follow how an incentivized revealed-preference measure and cognitive test scores predict the same health behavior in the same adults over twelve years, across a familiar and a new vaccine. Revealed-preference consistency and related measures of decision quality have mostly been studied where people largely agree on what a better choice is. Consistency predicts household wealth (Choi et al., 2014), a broader measure of decision-making ability predicts costly borrowing (Carvalho et al., 2024), and the quality of health-plan choices is judged against dominance or expected cost (Abaluck and Gruber, 2011; Bhargava et al., 2017). The closest evidence for health finds no statistically significant association between consistency and past-year flu vaccination, conditional on IQ and other controls (Li et al., 2018). The average association is close to zero here as well, but consistency predicts where receipt falls over time. Health behaviors are different, because people value health risks, side effects, and effort differently. Cognitive ability, by contrast, is a well-documented predictor of health. Higher test scores in youth predict lower mortality decades later (Calvin et al., 2011), cognitive ability accounts for a sizable share of the education gradient in smoking, obesity, and heavy drinking (Cutler and Lleras-Muney, 2010), and numeracy affects how people understand risk information and

make medical decisions (Peters et al., 2006; Reyna et al., 2009). In the same adults, the two measures carry different information: consistency predicts the decline in a routine vaccine, and quantitative ability predicts how quickly adults took up a new one.

Second, the paper provides suggestive evidence that beliefs work differently for a familiar and a new vaccine. Perceived risks of illness and of side effects are among the strongest correlates of vaccination (Brewer et al., 2007; Schmid et al., 2017), and these beliefs account for much of the variation in influenza demand and respond to information (Sacks and Sydnor, 2026; Sacks and Walker, 2026). To show that a gap in beliefs comes before a gap in receipt, we need beliefs measured before people can get a dose, and the literature says little about whether people with the same beliefs differ in how often they make mistakes when they decide. For COVID-19, the quantitative-ability gap in favorable safety and effectiveness beliefs was already present when 1.9 percent of the sample had a dose, did not widen during the rollout, and accounts for about 60 percent of the later ability gap in receipt. Relative to studies that explain vaccine demand through beliefs and preferences (Sacks and Sydnor, 2026), the results suggest that beliefs matter in different ways for the two vaccines. For a new vaccine, what matters is the beliefs people form before they can get a dose. For a familiar one, it is how often people make mistakes in a repeated decision once the perceived benefit falls.

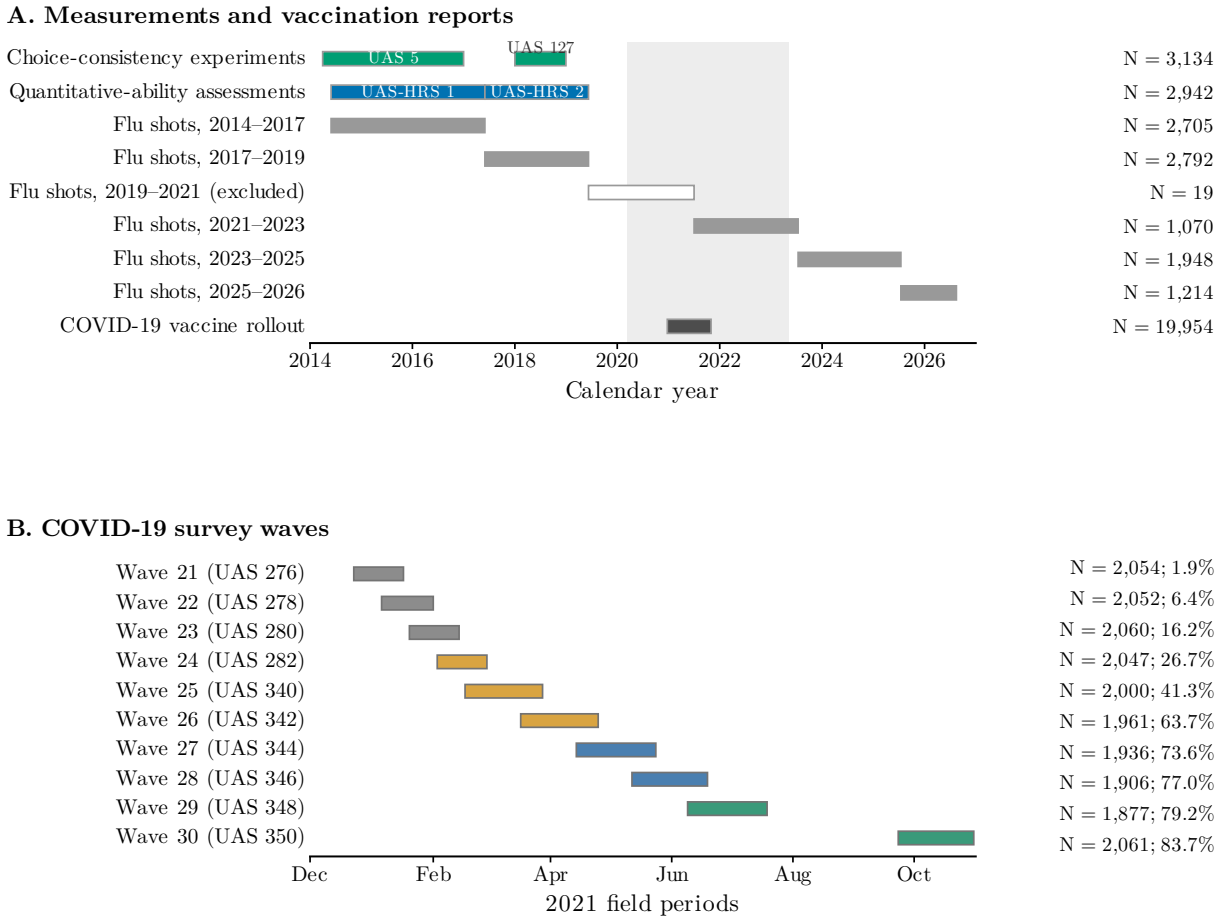
2 Data and the Two Measures

2.1 The Understanding America Study

The UAS is a probability-based internet panel that recruits U.S. adults through address-based sampling and provides internet access and tablets to households that could not otherwise participate online (Alattar et al., 2018). I link four UAS sources by respondent (Appendix A). The risky-budget experiments in Modules 5 and 127 give choice consistency. The UAS Comprehensive Panel Dataset harmonizes UAS–HRS Waves 1–6 and gives the cognitive assessments, the influenza vaccination reports, and demographic and education variables. UAS 71 records party identification, and the UAS COVID-19 National Sample Longitudinal File records vaccination status and vaccine beliefs in each wave of the 2021 rollout. Figure 1 places the measurements and outcomes on a common calendar.

The influenza analysis uses 9,729 respondent-period observations from 3,038 respondents with the choice-consistency measure. Requiring quantitative ability leaves 9,629 observations from 2,939 respondents. The COVID-19 regressions use a common sample of 19,954 resolved respondent-wave observations from 2,364 respondents, all of whom also appear in

Figure 1: Timing of Measurements and Analysis Samples



Notes: Data are from the Understanding America Study (UAS). Panel A shows when each measure and outcome was collected: the choice-consistency experiments (UAS 5 in 2014–2016 and UAS 127 in 2018), the two UAS–HRS waves that supply the quantitative-ability scores, the six influenza reporting periods, and the 2021 COVID-19 rollout. Gray shading marks the U.S. public health emergency (March 2020–May 2023). The 2019–2021 period is excluded from all regressions because only 19 linked respondents have a valid flu-shot report. Counts at the right are respondents for the two measures, respondent-period observations for the influenza periods, and respondent-wave observations for the rollout. Panel B shows the field period of each COVID-19 wave, with the UAS module number in parentheses, the number of respondents in the joint sample, and the unweighted share who had received at least one dose. Bar colors mark the initial rollout (Waves 21–23, gray), eligibility expansion (24–26, orange), broad adult availability (27–28, blue), and late rollout (29–30, green). Field periods of adjacent waves overlap.

the influenza analysis (Appendix Table A1). Respondents in the influenza sample averaged 62 years of age in 2021, and 76.5 percent are non-Hispanic White (Table 1, Panel D). So the results describe an older population, in which influenza vaccination is common.

2.2 Choice Consistency

Choice consistency comes from incentivized graphical budget experiments adapted from Choi et al. (2007, 2014). Module 5 fielded them in 2014–2016 and Module 127 in 2018. No respondent completed both, so each of the 1,461 respondents in Module 5 and the 1,673 in Module 127 contributes one measurement. In each of 25 scored rounds, the screen showed a budget line whose axes were dollar payoffs in two equally likely states, Heads and Tails, and the respondent chose a point on the line (Appendix Figure B1). Slopes varied across rounds, so moving along a line traded payoffs across states at different relative prices. In every round, at least one intercept was \$50 or more, and neither exceeded \$100. After the 25 choices, the program picked one round at random and drew Heads or Tails with equal probability, and the respondent received the allocation for the realized state. No outcome feedback was given between rounds.

The first index, e^*, is Afriat's critical cost efficiency index (CCEI). The Generalized Axiom of Revealed Preference (GARP) rules out revealed-preference cycles that contain a strict comparison, and choices satisfy GARP if and only if a continuous, increasing, and concave utility function rationalizes them (Afriat, 1967; Varian, 1982). The CCEI is the supremum of the budget-retention factors $e \in [0, 1]$ for which the choices satisfy GARP after every budget is scaled by e (Afriat, 1973; Choi et al., 2014). A value of 0.80 means the choices become consistent once every budget shrinks by 20 percent. It is not the share of choices answered correctly. One minus the CCEI can be interpreted as an upper bound on the fraction of the budget wasted through inconsistent choices, not as the fraction actually lost or the number of mistakes (Choi et al., 2014). The second index, e^{**}, applies the same algorithm after adding the mirror image of each observation, which swaps both state prices and state allocations. Consistency of the augmented data also requires that preferences be symmetric across the two equally likely states and monotone with respect to first-order stochastic dominance (Nishimura et al., 2017; Polisson et al., 2020). For example, a 50–50 chance of \$20 or \$40 is the same lottery whichever state pays \$20, and a 50–50 chance of \$30 or \$50 dominates it. Because the augmented data contain the original observations, $0 \leq e^{**} \leq e^* \leq 1$ for every respondent. The augmented criterion also treats violations of first-order stochastic dominance as mistakes, so e^{**} applies a stricter standard than e^* rather than counting more choices as mistakes. Neither index identifies a respondent's risk

preferences, and respondents with very different attitudes toward risk can attain the same value. I call the concept choice consistency, or consistency for short, and focus on e^{**}, which imposes the stronger standard. Appendix B gives formal definitions.

Of the 3,134 respondents, 3,066 have all 25 usable choices and 68 have between 2 and 24. Fewer rounds impose fewer revealed-preference restrictions, so they cannot lower the index for a given choice record. The main analysis keeps all 3,134 respondents, and Section 6.2 restricts the sample to the 3,066 with all 25 rounds.

2.3 Quantitative Ability

The quantitative-ability index averages four UAS assessments of numerical reasoning (Alattar et al., 2018; Gatz et al., 2026). Numeracy measures the ability to understand and use probabilities and other numerical information (Weller et al., 2013). Number series measures fluid reasoning in a quantitative task. Serial sevens measures mental arithmetic, sustained attention, and working memory (Gatz et al., 2026). Financial literacy measures knowledge and calculations used in financial decisions (Lusardi and Mitchell, 2014; Angrisani and Kapteyn, 2024). Numeracy and number series are item-response T-scores, serial sevens ranges from 0 to 5, and financial literacy from 0 to 14. Appendix B describes the items and scoring.

The four assessments were not always given in the same interview, so I use each respondent's first nonmissing score on each. I exclude from every analysis the 24 of the 3,158 linked experiment respondents whose first score on at least one assessment comes from UAS–HRS Wave 3 or Wave 4, so no measure in the paper was first observed after March 2020. Of the remaining 3,134, 2,942 have all four scores from Wave 1 or Wave 2 and form the construction sample. The other 192 lack at least one score, most often number series, and are left out of every model that includes quantitative ability. I standardize each assessment in the construction sample, average the four with equal weights, and standardize the average in the same sample.

2.4 What the Two Measures Capture

Quantitative ability and consistency capture different decision-making abilities. Quantitative ability reflects how well people understand and work with numerical information, so it may affect the beliefs they form about health risks. Consistency reflects how few mistakes people make when they make the same kind of decision again and again. Making consistent choices across 25 budgets requires paying attention to each one and thinking about which option is best for oneself. Many health behaviors ask the same of people, from a flu shot every season to a healthy diet every day. In Dutch panel data, consistency predicts lower rates of smoking

and obesity, outcomes built from many such choices (Li et al., 2018). A respondent who favors vaccination but makes more mistakes may likewise skip it in some seasons.

Quantitative ability behaves like a conventional cognitive measure: it correlates 0.46 with years of education, 0.49 with picture vocabulary, and 0.56 with verbal analogies (Appendix Table B2). It is also measured reliably. Three of its four components, numeracy, number series, and financial literacy, have correlations of 0.74–0.77 between UAS–HRS Waves 1 and 2 (Appendix Table B4). Serial sevens is the exception, at 0.31. Its low correlation does not mean that most respondents' performance changed: many score five in both waves, so the few who move between waves drive the correlation. The ceiling mainly limits how well this test separates respondents.

Most respondents choose more consistently than random choice would. Depending on the module and index, 57–62 percent of observed values exceed the 95th percentile of the distribution generated by random choices on the same budgets (Bronars, 1987; Figure 2). Not exceeding this threshold does not establish random choice or inattention. In the analysis sample, e^* averages 0.875 and e^{**} averages 0.750, with standard deviations of 0.149 and 0.231 (Table 1). So a one-standard-deviation difference in e^{**} is equivalent to a critical budget contraction 23.1 percentage points smaller. The module means are 0.879 in UAS 5 and 0.873 in UAS 127 for e^*, and 0.763 and 0.738 for e^{**}. About 21 percent of respondents have e^* equal to one, against 2.4 percent for e^{**} (Appendix Table B1).

The two kinds of measure rank people differently. Each consistency index correlates 0.27 with quantitative ability, and the two indices correlate 0.75 with each other (Appendix Table B2). The consistency indices correlate 0.09–0.27 with the four quantitative-ability components, years of education, picture vocabulary, and verbal analogies. Consistency is higher among younger respondents, men, and those with more education and income, but these characteristics together explain only 3.7–4.4 percent of its variation (Appendix Table B3). Choi et al. (2014) report a similar pattern for a Dutch panel.

The choice experiment was administered once, so its test-retest reliability is unknown. Splitting each respondent's 25 rounds into odd and even halves gives split-half correlations of 0.41 for e^* and 0.59 for e^{**}, and the Spearman–Brown formula, which adjusts for using half the rounds, puts the reliability of the full indices at 0.58 and 0.75 (Appendix Table B4). Each half contains only 12 or 13 rounds, which weakens the revealed-preference restrictions, so these figures may understate the reliability of the full indices. They also come from a single sitting, so they do not capture changes from one occasion to the next.

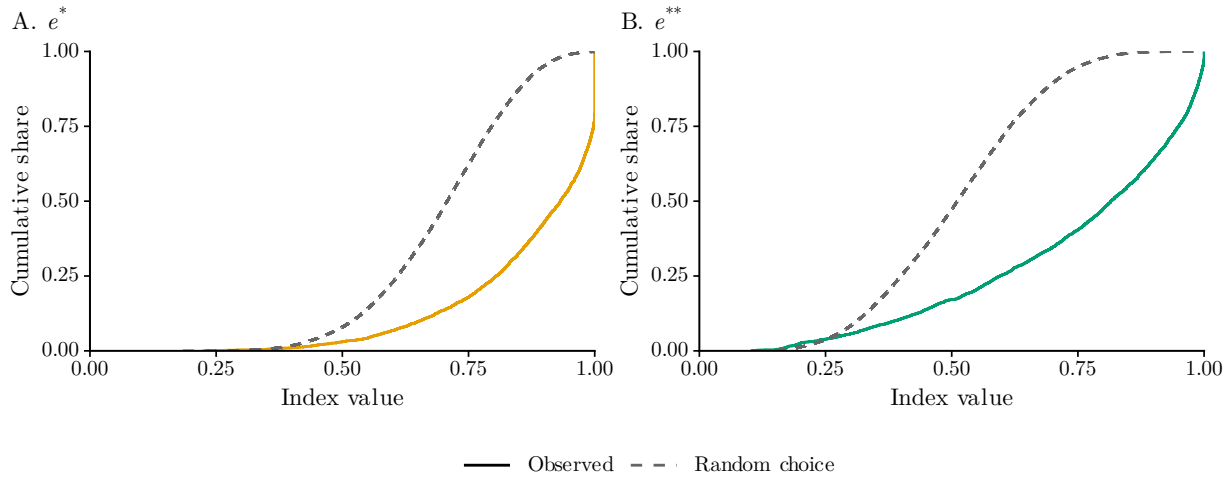
Part of the low end of the consistency distribution reflects how carefully respondents completed the experiment. Following Leiner (2019), I divide the median time on each scored, instruction, and practice screen, within module and device type, by the respondent's time, cap

Table 1: Summary Statistics

	(1) Observations	(2) Mean	(3) SD	(4) Min	(5) Max
Panel A. Choice consistency and quantitative ability					
e^*	3,134	0.88	0.15	0.22	1.00
e^{**}	3,134	0.75	0.23	0.10	1.00
Quantitative ability (index)	2,942	−0.00	1.00	−4.00	1.90
Numeracy (T-score)	3,134	50.85	8.58	33.54	70.45
Number series (T-score)	2,951	50.88	9.11	13.64	65.27
Serial sevens (0–5)	3,042	4.57	0.98	0.00	5.00
Financial literacy (0–14)	3,130	9.93	3.04	0.00	14.00
Panel B. Influenza vaccination					
2014–2017	2,705	0.623	0.485		
2017–2019	2,792	0.600	0.490		
2021–2023	1,070	0.650	0.477		
2023–2025	1,948	0.616	0.487		
2025–2026	1,214	0.570	0.495		
Panel C. COVID-19 receipt and Wave 21 views					
Cumulative receipt: initial rollout	6,166	0.082	0.274		
Cumulative receipt: eligibility expansion	6,008	0.437	0.496		
Cumulative receipt: broad availability	3,842	0.753	0.431		
Cumulative receipt: late rollout	3,938	0.816	0.388		
Favorable effectiveness belief, Wave 21	2,053	0.74	0.26	0.00	1.00
Favorable safety belief, Wave 21	2,052	0.79	0.24	0.00	1.00
CDC/HHS trust, Wave 21 (1–4)	2,048	2.48	0.83	1.00	4.00
Panel D. Baseline characteristics					
Female	2,939	0.526	0.499		
Age in 2021	2,939	62.22	13.14	25.00	103.00
Asian	2,939	0.027	0.162		
Hispanic	2,939	0.065	0.247		
Non–Hispanic Black	2,939	0.082	0.275		
Non–Hispanic White	2,939	0.765	0.424		
Other race	2,939	0.060	0.238		
Below high school	2,939	0.058	0.233		
Associate degree	2,939	0.138	0.345		
Bachelor's degree	2,939	0.216	0.412		
Graduate degree	2,939	0.165	0.372		
High school	2,939	0.180	0.385		
Some college, no degree	2,939	0.243	0.429		

Notes: Each panel uses a different unit of observation. Means of binary variables are proportions. Panel A has one observation per respondent. The sample is the 3,134 choice-experiment respondents who remain after the pre-pandemic timing restriction; some rows have fewer observations because some respondents lack that score. The quantitative-ability index combines four component scores. It is defined only for respondents with all four components by UAS–HRS Wave 2 and is standardized to mean 0 and standard deviation 1 among those respondents. Panel B has one observation per respondent and reporting period. The mean in each row is the share of observations in that period with a reported flu shot. Panel C has one observation per respondent and survey wave. Cumulative receipt is the share of observations in each rollout phase in which the respondent has reported at least one dose. The three Wave 21 rows use only Wave 21 respondents. Favorable effectiveness belief equals $(100-p)/100$, where p is the stated probability (0–100) that a vaccinated person will contract COVID-19. Favorable safety belief applies the same formula to the stated probability that a vaccine will cause serious side effects. Both beliefs range from 0 to 1, with higher values meaning more favorable beliefs. Institutional trust is the mean of two items (trust in the CDC and HHS), each scored from 1 to 4, among respondents with valid answers to both items; higher values mean more trust. Panel D describes the 2,939 respondents in the influenza sample with at least one usable influenza report and complete choice-consistency, quantitative-ability, and baseline covariate data.

Figure 2: Choice Consistency Relative to Random Choice



Notes: Data are from the UAS choice experiments in Modules 5 and 127. Each panel plots the cumulative distribution of the index for the 3,134 respondents and the distribution produced by random choice on the same budgets (Bronars, 1987). The random-choice distribution is simulated 10,000 times per module, drawing each allocation uniformly along the budget line. Depending on the module and index, 57–62 percent of observed values exceed the module-specific 95th percentile of the random-choice distribution.

the ratio at three, and average across screens. An average of at least two flags fast responses. An average below one-half flags slow responses, since inattention need not be limited to fast respondents (Read et al., 2022). Among influenza respondents with timing data, 6.2 percent are fast and 3.2 percent slow (Appendix Table B1, Panel C). Fast respondents fail the experiment's final paid interface task more often than other respondents with an observed result, 47.0 against 15.2 percent, and their mean e^{**} is 0.60, against 0.76 among respondents with typical speed. Speed and interface performance are not perfect indicators of attention, so we should not read fast or slow respondents as purely inattentive. Section 6.2 reports how the results change when fast respondents are excluded, at this cutoff and at a stricter cutoff of 1.5.

3 Conceptual Framework

A simple model of the vaccination decision explains why a measure of how often people make mistakes in repeated decisions can have a near-zero average association with vaccination and still predict who changes course. It also explains why skill with new information should matter most when a vaccine is new. In the model, choice consistency and quantitative ability stand in for two kinds of precision. I do not estimate the model. I use it to state predictions that the vaccination outcomes could contradict (Appendix Table C1).

3.1 The Decision Rule

Consider individual i, vaccine v, and decision date t. Over a fixed health horizon, let $p^v_{0,it}$ and $p^v_{1,it}$ be the perceived probabilities of illness without and with vaccination, $L^v_i > 0$ the utility loss from illness, r^v_{it} the perceived probability of side effects, $S^v_i > 0$ their utility loss, and $k^v_{it} \geq 0$ the anticipated monetary, time, and effort cost of getting a dose. The perceived net gain from vaccination is

$$G^v_{it} = \left(p^v_{0,it} - p^v_{1,it}\right) L^v_i - r^v_{it} S^v_i - k^v_{it}. \tag{1}$$

The first term is the expected health benefit, the second the expected cost of side effects, and the third the cost of getting a dose. Let $F^v_{it} \in \{0, 1\}$ indicate that vaccination is feasible and u^v_{it} be an evaluation disturbance in the same units. Receipt is

$$a^v_{it} = F^v_{it}\, \mathbf{1}\left\{G^v_{it} + u^v_{it} \geq 0\right\}. \tag{2}$$

So a person gets vaccinated when a dose is available and her overall assessment, the gain plus the disturbance, is not negative. The disturbance captures decision-specific variation in how she combines the parts of G^v_{it} into one assessment. Systematic preferences, including any value placed on protecting others, are part of G^v_{it}, so the precision parameter introduced next does not absorb them.

3.2 Decision Precision and a Familiar Vaccine

For a familiar vaccine, what matters is how often people make mistakes when they decide on the basis of assessments they already have. Beliefs about personal benefits and side effects account for much of the variation in influenza demand and respond to information (Sacks and Sydnor, 2026; Sacks and Walker, 2026). I take those assessments as given and ask how reliably people choose in line with them. Let $c_i > 0$ be decision precision and $\sigma_t > 0$ a common, period-specific scale of decision difficulty, and let $u^F_{it} = (\sigma_t / c_i)\varepsilon_{it}$, with ε_{it} standard normal and independent of the gain. In random-utility accounts of the budget experiment, noise in evaluating allocations produces revealed-preference cycles, so the CCEI falls as evaluation noise rises (Choi et al., 2014; Apesteguia and Ballester, 2015; Dean and Martin, 2016). So I use consistency as a proxy for c_i. Conditional on feasibility, the probability of influenza vaccination and its precision gradient are

$$P^F_{it} = \Phi(z_{it}), \qquad \frac{\partial P^F_{it}}{\partial c_i} = \frac{z_{it}}{c_i}\,\phi(z_{it}), \qquad z_{it} = \frac{c_i G^F_{it}}{\sigma_t}, \tag{3}$$

where Φ and ϕ are the standard normal distribution and density functions. A mistake is a choice that the gain does not support: declining a dose when $G_{it}^F > 0$, or taking one when $G_{it}^F < 0$. Conditional on feasibility and a nonzero gain, its probability in a given decision, $\Phi(-c_i|G_{it}^F|/\sigma_t)$, falls with precision and is largest when the gain is close to zero. Two implications follow (Appendix C.1).

Implication F1. Greater decision precision raises the probability of vaccination when the perceived net gain is positive and lowers it when the gain is negative. The average gradient is therefore near zero when gains of opposite signs offset one another, or when most assessments are settled far from indifference, where precision does not affect the choice.

Implication F2. Consider two respondents with the same positive gain who are both likely to vaccinate, with probabilities above $\Phi(1) \approx 0.84$. A less favorable environment, meaning a lower perceived gain or greater decision difficulty, reduces vaccination more for the less precise respondent. The precision gradient then widens through a decline at the bottom rather than a rise at the top. A more favorable environment produces the reverse pattern.

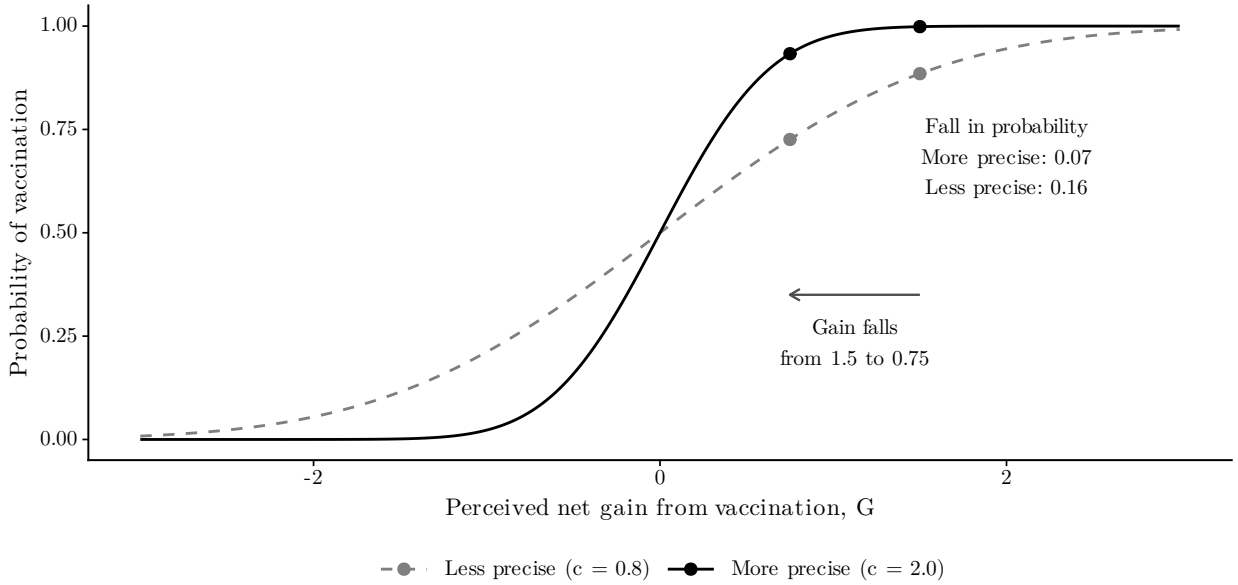
In words, when a familiar vaccine starts to look less worthwhile, vaccination falls more among people who make more mistakes in deciding. Figure 3 illustrates both implications. The two curves cross at indifference, so precision can have a near-zero average effect when gains of opposite signs offset one another. When the gain falls from 1.5 to 0.75, the probability of vaccination falls by 0.07 for the more precise respondent and by 0.16 for the less precise one. The same logic orders groups by risk. Respondents who perceive a high risk of illness are likely to have gains far above zero, so a common deterioration changes their vaccination little at any level of precision. Under these conditions, the widening can be larger among respondents with lower perceived risk, whose gains are positive but closer to zero, as long as both kinds of respondents remain likely to vaccinate.

The influenza outcome records any receipt over about two seasonal decisions. With independent seasonal disturbances, both implications carry over to this two-season outcome (Appendix C.1). The framework does not say which change occurs: lower perceived protection, greater side-effect concern, higher practical cost, and more conflicting information all lower the gain or raise the difficulty.

3.3 Information, Access, and a New Vaccine

People have little experience with a new vaccine, so they revise their assessments as evidence comes in, and whether a revised assessment leads to receipt depends on access. Suppose that individual i processes common evidence with precision η_i and accepts vaccination once her

Figure 3: Decision Precision and the Probability of Vaccination



Notes: The curves plot $P = \Phi(cG/\sigma)$ from Equation (3) with $\sigma = 1$, for a more precise ($c = 2.0$) and a less precise ($c = 0.8$) respondent. The marked points show the probability of vaccination before and after the perceived gain falls from 1.5 to 0.75. The parameter values are illustrative and are not estimated.

updated assessment of the net health benefit exceeds the cost of getting a dose. Along a common favorable path of evidence, the number of signals she needs before accepting falls with η_i (Appendix C.2). Let T_i^{info} be the date she accepts and $T_i^{\text{feas}} = T_i^{\text{elig}} + \Delta_i$ the first date a dose is feasible, where T_i^{elig} is the date she becomes eligible, set by age, occupation, health, and state of residence, and $\Delta_i \geq 0$ is the time she needs to get a dose once eligible. Once access opens it stays open, so first receipt occurs at

$$T_i = \max\left\{T_i^{\text{info}},\, T_i^{\text{feas}}\right\}. \tag{4}$$

In words, a person gets vaccinated once she is both willing and able.

I use quantitative ability as a proxy for η_i. Finding and booking an appointment is itself a quantitative task, so I also use it as a proxy for a shorter Δ_i. Numeracy predicts the use of numerical risk information and resistance to misinformation (Peters et al., 2006; Cokely et al., 2012; Roozenbeek et al., 2020). In a UAS survey experiment, respondents with lower numeracy viewed positive and negative influenza-vaccine narratives as more informative and changed their judged vaccination probabilities more (Bruine de Bruin et al., 2017).

Implication C1. If information binds, higher-ability respondents revise their beliefs favorably faster and receive a dose weakly earlier.

Implication C2. If access or arranging binds, a receipt gradient by ability opens when

eligibility opens, even among respondents who accepted vaccination on the same date, and narrows when getting a dose becomes easier, for example through pre-booked appointments.

Implication C3. Once both higher- and lower-ability respondents have accepted vaccination and access is open, equal subsequent hazards shrink an existing gap in cumulative receipt without reversing it (Appendix C.3). So a narrowing gap does not by itself show that differences in learning have diminished.

Learning can also happen before doses are available, because evidence about a new vaccine starts to accumulate before the rollout. If higher-ability respondents update faster in that period, they enter the rollout with more favorable beliefs, and the belief gap need not widen afterward. By Equation (4), their receipt gap still opens only when access opens.

4 Vaccination Outcomes, Samples, and Specifications

4.1 Influenza Vaccination

The influenza outcome is self-reported receipt of a flu shot during a recall interval of about two years, from UAS–HRS Waves 1, 2, 4, 5, and 6. The first UAS–HRS interview asks about a flu shot in the previous two years. Each follow-up interview asks about receipt since the respondent's preceding UAS–HRS interview, and respondents who join the series later get the first-interview wording. Yes is coded one and no zero; don't-know, refusal, item-missing, and routed-missing responses are treated as missing. An interval of about two years spans two influenza seasons, so the outcome records whether the respondent got at least one dose across roughly two seasonal decisions. It is not an annual vaccination rate, and it is not verified against records. Wave 3 (2019–2021) is omitted because most respondents with a valid choice-consistency measure lack a valid answer to the flu-shot question, which leaves only 19 usable records.

Reporting periods are labeled by the years in which their interviews took place: 2014–2017, 2017–2019, 2021–2023, 2023–2025, and 2025–2026 (Figure 1, Panel A). The 2025–2026 interviews in the September 2026 data release run through August 16, 2026 and cover roughly the 2023–24 through 2025–26 seasons. The recall interval is 24 months for a first interview and the time since the preceding interview otherwise, so the 2021–2023 interval begins at the Wave 3 interview even though Wave 3 supplies no usable report. Two restrictions limit overlap between periods. I drop 98 follow-up records whose immediately preceding source wave is missing; the remaining intervals average 24–25 months. For 2021–2023, I also drop the 1,110 interviews conducted before May 1, 2022, whose intervals are especially likely to include months before March 2020. Some remaining intervals still exceed 24 months, so this

restriction reduces but does not remove pre-pandemic overlap, and later interviewees may differ from earlier ones.

The five periods contain 2,705; 2,792; 1,070; 1,948; and 1,214 observations, and a flu shot is reported in 61.1 percent of the 9,729 observations. Survey and observation weights leave the UAS time pattern broadly unchanged. Age standardization lowers the UAS rates and narrows their gap with the National Health Interview Survey, so the sample's older age composition accounts for part of the unweighted difference (Appendix Figure A1).

4.2 COVID-19 Vaccination During the Rollout

The COVID-19 outcome is cumulative self-reported receipt of at least one dose by each interview in Waves 21–30 of the UAS COVID-19 National Sample Longitudinal File, fielded from December 23, 2020 through October 31, 2021. The indicator switches to one when a direct response, or confirmation of a preloaded earlier report, first establishes receipt. It stays one in every later wave in which the respondent participates, even if a later response is missing or inconsistent. Before receipt is established, a direct no is coded zero; unsure, routed-missing, and unresolved responses are treated as missing. The records do not give complete, internally consistent first-dose dates, so I index vaccination status by interview date.

I group the waves into four rollout phases (Figure 1, Panel B): initial rollout (Waves 21–23), eligibility expansion (Waves 24–26), broad adult availability (Waves 27–28), and late rollout (Waves 29–30). The phases follow the survey field periods and the national move from priority-group allocation to nationwide adult eligibility by April 19, 2021 (Diesel et al., 2021). Adjacent field periods overlap, so different waves can include interviews on the same date. The phases do not measure individual eligibility, which depended on age, occupation, residence in long-term care, health conditions, and state.

Each wave also asks about two vaccine beliefs on a 0–100 probability scale: the probability that a vaccinated person will still contract COVID-19, and the probability that a vaccine will cause serious side effects. I code favorable effectiveness beliefs as 100 minus the first and favorable safety beliefs as 100 minus the second. They correspond to the perceived protection and side-effect risk in Equation (1).

4.3 Controls and Additional Characteristics

The main specifications control for fixed baseline characteristics: sex; birth year, entered linearly; race and ethnicity in five categories (Hispanic of any race, and non-Hispanic White, Black, Asian, or other race, the last including American Indian or Alaska Native, Native

Hawaiian or Pacific Islander, and multiracial respondents); and educational attainment in six categories (below high school, high school, some college without a degree, associate degree, bachelor's degree, and graduate degree). Education is the category reported in UAS–HRS Wave 1, or in Wave 2 if Wave 1 is missing, and is held fixed. I leave income, health, and insurance out of the main models because they may themselves reflect quantitative ability or consistency, so conditioning on them could absorb part of the associations of interest. Table 1 reports summary statistics.

The channel analyses in Section 5 and the robustness checks in Section 6 add further characteristics, each with period- or phase-specific coefficients. Party identification comes from UAS 71 (November 9–December 19, 2016), grouped as Democrat, Republican, independent, or other. An institutional-trust index averages Wave 21 trust in the CDC and HHS, each coded from one (do not trust at all) to four (trust completely), and is standardized using its mean and standard deviation among all Wave 21 respondents with valid answers to both items. Household income (inverse hyperbolic sine), self-rated health, a count of seven chronic conditions, and health-insurance coverage take their first nonmissing UAS–HRS value in Waves 1–2. Risk attitudes are measured by the mean share of the experimental budget allocated to the cheaper state, standardized within module. Personality is measured by the five official UAS Big Five scores, each from Wave 1 or, if missing, Wave 2, all measured before March 2020. Response speed follows Leiner (2019), as described in Section 2.4, and a separate restriction keeps respondents who passed the experiment's final paid interface task under its bonus-payment rule.

4.4 Specifications

The coefficients of interest are changes in conditional associations over time. Stable omitted differences can shift the level of an association but not its change. Changes in sample composition, or in how other characteristics relate to vaccination, can shift the change itself. If the education gradient in vaccination changed after 2020, for example, a fixed education coefficient would push part of that change into the consistency interaction, because education and consistency are correlated. So the models let the demographic and education coefficients vary with the time comparison. All main models are unweighted linear probability models with standard errors clustered by respondent; Section 6.4 reports weighted checks (Solon et al., 2015). Consistency enters on its zero-to-one scale and quantitative ability in construction-sample standard deviations. Tables, figures, and text report associations in percentage points per standard deviation, using the consistency index's standard deviation across unique respondents in the estimation sample and the fixed construction-sample units

for quantitative ability.

I estimate three influenza models, which divide time more finely in turn. The average-association model is

$$\text{FluShot}_{it} = \lambda_t + Z_{ij}'\beta_j + X_i'\gamma + \rho\,\text{Recall}_{it} + u_{it}, \tag{5}$$

where i indexes respondents and t reporting periods. For $j \in \{*, **\}$, e_i^j is one of the two consistency indices and A_i is quantitative ability, and the vector Z_{ij} contains e_i^j, A_i, or $(e_i^j, A_i)'$, depending on whether a specification includes one measure or both. The vector X_i contains the baseline controls, Recall_{it} is the number of months since the preceding interview, and the reporting-period fixed effects λ_t absorb differences in average vaccination across periods. The vector β_j gives the association averaged across the five periods (Table 2, Panel A).

The second model compares the two reporting periods before 2020 with the three after:

$$\text{FluShot}_{it} = \lambda_t + Z_{ij}'\beta_j + \text{Post}_t\, Z_{ij}'\delta_j + X_i'\gamma + \text{Post}_t\, X_i'\kappa + \rho\,\text{Recall}_{it} + u_{it}. \tag{6}$$

Here Post_t equals zero in 2014–2017 and 2017–2019 and one in 2021–2023, 2023–2025, and 2025–2026, and its main effect is absorbed by λ_t. The vector β_j gives the association in 2014–2019 and δ_j its change in 2021–2026, so the association in the later periods is $\beta_j + \delta_j$. The vector κ gives the corresponding changes in the coefficients on X_i (Table 2, Panel B).

The third model replaces the single Post_t interaction with separate reporting-period interactions:

$$\text{FluShot}_{it} = \lambda_t + Z_{ij}'\beta_{j0} + \sum_{\tau\neq 0} \mathbf{1}\{t=\tau\}\, Z_{ij}'\beta_{j\tau} + X_i'\gamma + \sum_{\tau\neq 0} \mathbf{1}\{t=\tau\}\, X_i'\kappa_\tau + \rho\,\text{Recall}_{it} + u_{it}, \tag{7}$$

where $\tau \in \{0,1,2,3,4\}$ indexes the five reporting periods in chronological order and $\tau = 0$ is 2014–2017. The vectors β_{j0} and γ give the coefficients on Z_{ij} and X_i in 2014–2017, and $\beta_{j\tau}$ and κ_τ give their changes in period τ, so the association in period τ is $\beta_{j0} + \beta_{j\tau}$ (Appendix Table D1).

For COVID-19, I first estimate how the associations vary across the four rollout phases:

$$\begin{aligned} C_{iw} = \lambda_w + Z_{ij}'\beta_j + \sum_{q\neq 0} D_{wq}\, Z_{ij}'\delta_{jq} + X_i'\gamma + \sum_{q\neq 0} D_{wq}\, X_i'\kappa_q \\ + \text{State}_i'\psi_0 + \sum_{q\neq 0} D_{wq}\,\text{State}_i'\psi_q + \sum_{\omega} \rho_\omega\, d_{iw}\, \mathbf{1}\{w=\omega\} + u_{iw}. \end{aligned} \tag{8}$$

Here w indexes survey waves and C_{iw} is the cumulative receipt indicator. The index $q \in \{0, 1, 2, 3\}$ denotes rollout phase, with $q = 0$ the initial rollout, and D_{wq} equals one when wave w belongs to phase q. Wave fixed effects λ_w absorb separate phase main effects. The vector State_i contains baseline-state indicators; ψ_0 gives state differences during the initial rollout, and ψ_q lets them change across later phases. The variable d_{iw} is the number of days between respondent i's interview and the mean interview date in wave w, and ρ_ω gives a separate within-wave interview-date slope for each wave ω. The vectors β_j and γ give the coefficients on Z_{ij} and X_i during the initial rollout, and each δ_{jq} gives the change in the association in phase q, so the phase-q association is $\beta_j + \delta_{jq}$ (Table 3). Cumulative receipt records vaccination through the interview date, so a phase change reflects differences built up through that phase, not vaccination within that phase alone.

To trace the timing more closely, I replace the phase interactions with separate interactions for each wave:

$$C_{iw} = \lambda_w + Z_{ij}'\beta_{j,21} + \sum_{\omega \neq 21} \mathbf{1}\{w = \omega\}\, Z_{ij}'\delta_{j\omega} + X_i'\gamma_{21} + \sum_{\omega \neq 21} \mathbf{1}\{w = \omega\}\, X_i'\kappa_\omega + \text{State}_i'\psi_{21} + \sum_{\omega \neq 21} \mathbf{1}\{w = \omega\}\, \text{State}_i'\psi_\omega + \sum_{\omega} \rho_\omega\, d_{iw}\, \mathbf{1}\{w = \omega\} + u_{iw}. \quad (9)$$

Wave 21 is the reference: $\beta_{j,21}$ gives the association in that wave and $\delta_{j\omega}$ its change in wave ω, so the association in wave ω is $\beta_{j,21} + \delta_{j\omega}$ (Appendix Table E1). The phase and wave interactions describe changes in conditional associations over calendar time, not effects of particular rollout policies.

Two estimates recur, and I name them. The flu divergence (e^{**}) is the change in the e^{**} association with influenza vaccination from 2014–2019 to 2021–2026: the element of δ_{**} on e^{**} in Equation (6), from the joint model with quantitative ability. The availability widening is the change in the quantitative-ability association with COVID-19 receipt from the initial rollout to broad adult availability: the element of $\delta_{**,q}$ on quantitative ability for that phase in Equation (8). Reported p-values are unadjusted.

5 Results

5.1 Influenza Vaccination

5.1.1 Small Average Associations That Grow over Time

Averaged across the five reporting periods between 2014 and 2026, neither consistency index predicts influenza vaccination, and the estimates rule out large associations. Conditional

Table 2: Choice Consistency, Quantitative Ability, and Influenza Vaccination

	(1) e^{**}	(2) Quantitative ability	(3) Both
Panel A. Average association, 2014–2026			
e^{**}	0.39		0.20
	(0.78)		(0.81)
Quantitative ability		1.20	1.16
		(0.93)	(0.95)
Observations	9,729	9,629	9,629
Respondents	3,038	2,939	2,939
Adjusted R^2	0.070	0.071	0.071
Panel B. Association in 2014–2019 and change in 2021–2026			
e^{**}: 2014–2019	−0.61		−0.71
	(0.83)		(0.86)
2021–2026 minus 2014–2019	2.25**		2.01**
	(0.93)		(0.95)
Quantitative ability: 2014–2019		0.28	0.44
		(1.02)	(1.04)
2021–2026 minus 2014–2019		2.26*	1.78
		(1.20)	(1.21)
Observations	9,729	9,629	9,629
Respondents	3,038	2,939	2,939
Mean outcome, 2014–2019 (%)	61.1	61.1	61.1
Mean outcome, 2021–2026 (%)	61.1	61.1	61.1
Adjusted R^2	0.072	0.072	0.072

Notes: Entries are associations with reported influenza vaccination, in percentage points per standard deviation of each measure, from unweighted linear probability models. Panel A estimates Equation (5), which pools the five reporting periods. Panel B estimates Equation (6) and reports the association in 2014–2019 and its change in 2021–2026. Columns 1 and 2 enter one measure at a time, and column 1 keeps the 99 respondents without a complete quantitative-ability index; column 3 enters both measures. All models include reporting-period fixed effects, months since the preceding interview, and baseline sex, birth year, race and ethnicity, and six education categories. In Panel B, the demographic and education coefficients differ before and after 2020; months since the preceding interview has a single coefficient. e^{**} is scaled by its standard deviation across unique respondents in the estimation sample, and quantitative ability by its standard deviation in the construction sample. Appendix Table F1 reports the same models with e^{*}. Standard errors clustered by respondent are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

on quantitative ability, a one-standard-deviation increase in e^{**} is associated with reported vaccination 0.2 percentage points higher (s.e. 0.81; Table 2, Panel A, column 3). The confidence intervals for e^{**} and e^{*} exclude positive associations above about 1.8 and 2.1 points, respectively (for e^{*}, Appendix Table F1, Panel A, column 2). Both bounds are about a fifth of the college gap: in the pre/post model with both measures, pre-2020 vaccination is 9.9 points higher among bachelor's-degree holders than among respondents without a high-school diploma (Appendix Table D2, Panel A, column 4). Quantitative ability has a larger point estimate, 1.2 points (s.e. 0.95), but its confidence interval runs from about -0.7 to 3.0 points.

The e^{**} association rises over the sample period, and the rise begins before the pandemic. Before 2020, the association is -0.7 points per standard deviation (s.e. 0.86; Table 2, Panel B, column 3). The flu divergence (e^{**}), its change from 2014–2019 to 2021–2026, is 2.01 points (s.e. 0.95; $p = 0.03$), about a fifth of the college gap. Relative to 2014–2017, the association is 1.7 points higher in 2017–2019 (s.e. 0.82), about 2.4 points higher in both 2021–2023 (s.e. 1.50) and 2023–2025 (s.e. 1.19), and 4.3 points higher in 2025–2026 (s.e. 1.47), more than two-fifths of the college gap (Appendix Table D1, column 3; Figure 4, Panel A). Because the rise starts before 2020, the average post-2020 association exceeds the 2017–2019 association by a smaller 1.3 points (s.e. 0.99), which is not statistically significant.

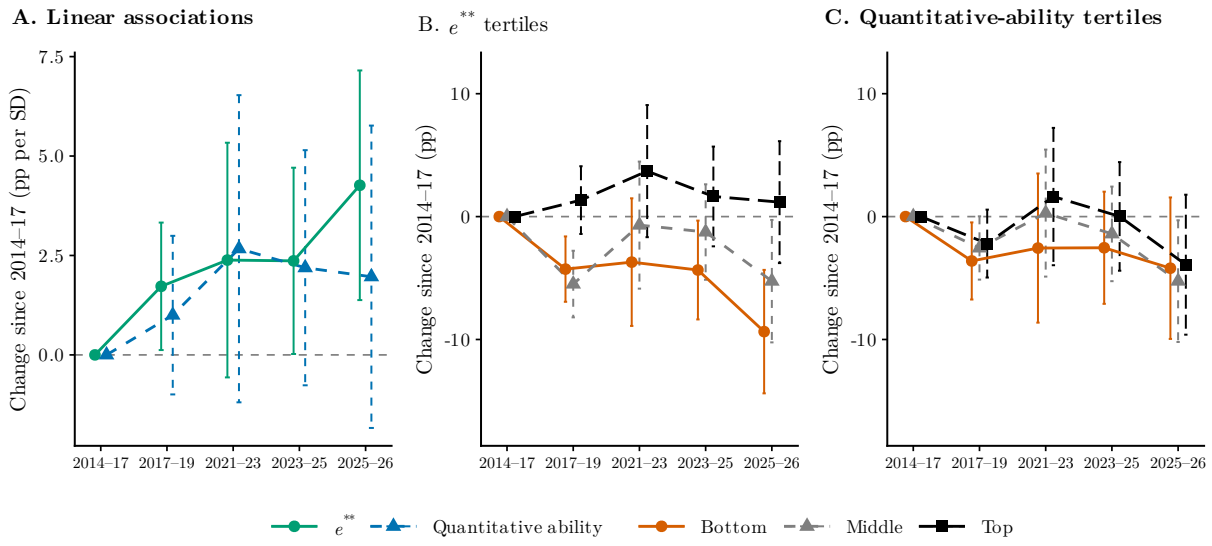
The changes for quantitative ability and e^{*} are similar in size to the flu divergence but not statistically significant. The quantitative-ability association rises by 1.78 points per standard deviation from 2014–2019 to 2021–2026 (s.e. 1.21; Table 2, Panel B, column 3). Unlike the e^{**} association, it does not keep rising: relative to 2014–2017, it is 2.7 points higher in 2021–2023 (s.e. 1.97) and 2.0 points higher in 2025–2026 (s.e. 1.94), when the e^{**} change is 4.3 points (Appendix Table D1, column 3). The e^{*} association rises by 1.62 points (s.e. 1.00; Appendix Table F1, Panel A, column 2).

5.1.2 The Divergence Comes from the Bottom of the Consistency Distribution

To locate the change, I replace each continuous measure with fixed tertiles interacted with reporting period and average the predictions over the same covariate distribution in every period (Figure 4, notes). The joint tertile model includes both measures and compares adjusted groups rather than following the same respondents.

Vaccination falls most in the bottom e^{**} tertile. Between 2014–2017 and 2025–2026, adjusted vaccination falls by 9.4 percentage points in the bottom tertile (s.e. 2.6), about the size of the college gap. In the top tertile, the estimated change is 1.2 points (s.e. 2.5; Figure 4, Panel B). The top-minus-bottom change is 10.5 points in 2025–2026 (s.e. 3.6), significant at the 1 percent level; the gap also widens by 7.4 points (s.e. 3.8) in 2021–2023

Figure 4: Influenza Vaccination by Reporting Period: Linear Associations and Tertile Paths



Notes: Panel A plots the change since 2014–2017 in the associations of e^{**} and quantitative ability with reported influenza vaccination, in percentage points per standard deviation, from the joint model in Equation (7). Panels B and C plot the covariate-adjusted change since 2014–2017 in the share reporting a flu shot, in percentage points, by fixed tertile of e^{**} and of quantitative ability. Both panels come from one linear probability model that interacts the e^{**} and quantitative-ability tertile indicators with reporting period. Sex, birth year, race and ethnicity, and education have period-specific coefficients, and months since the preceding interview has a single coefficient. Predictions are averaged over the covariates of the same 2,939 respondents in every period, with recall length set to its mean, and tertile cutoffs are fixed over time (Appendix Table D3). All panels use 9,629 observations, and the 2014–2017 values are zero by construction. Bars are pointwise 95 percent confidence intervals clustered by respondent.

and 6.0 points (s.e. 3.0) in 2023–2025 (Appendix Table D3, Panel A). The e^* tertiles follow the same broad path (Appendix Table F1, Panel C). Unadjusted rates by tertile show the same pattern (Appendix Figure D1).

The quantitative-ability tertiles decline more similarly. By 2025–2026, adjusted vaccination falls by 4.2 points in the bottom tertile (s.e. 2.9) and 3.9 points in the top tertile (s.e. 2.9); their 0.3-point difference is not statistically significant (Figure 4, Panel C; Appendix Table D3, Panel B).

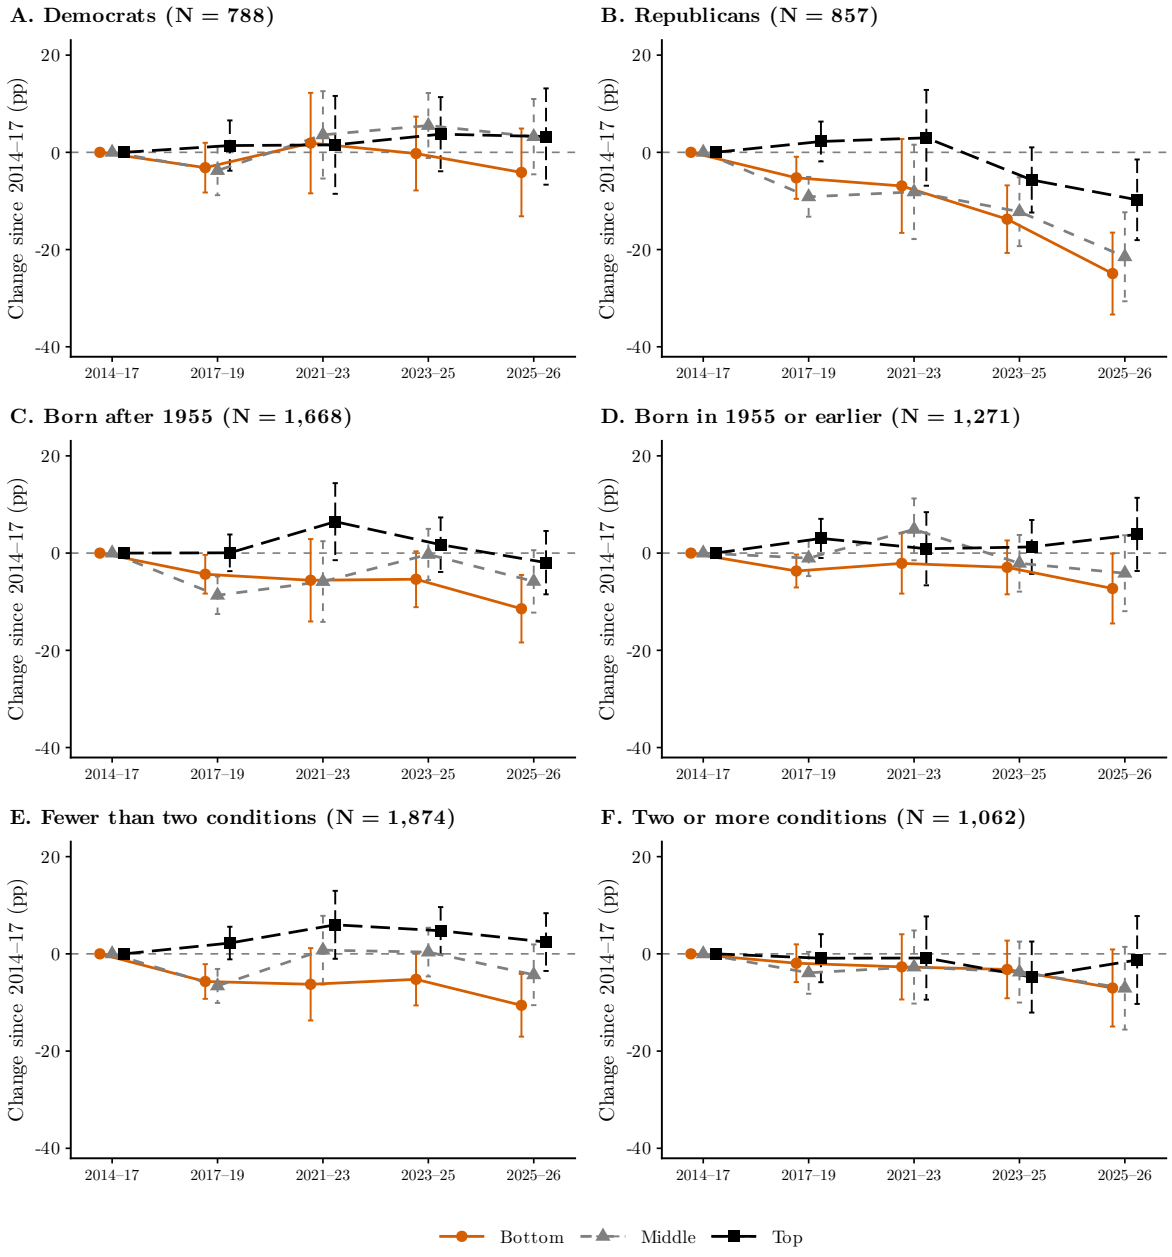
The separation between the e^{**} tertiles began in 2017–2019, when the top-minus-bottom change was already 5.6 points (s.e. 2.0). The 2017–18 season was unfavorable for influenza vaccination: interim vaccine effectiveness was 36 percent (Flannery et al., 2018), and national adult coverage fell from 43.3 to 37.1 percent (Centers for Disease Control and Prevention, 2018).

5.1.3 What Accounts for the Flu Divergence?

If the flu divergence reflects mistakes in a repeated decision, the fall in receipt among the least consistent respondents, largest in 2025–2026 (Section 5.1.2), should come mainly from those whose expected gain from a flu shot is small. A respondent with a large gain is likely to vaccinate even if her decisions are sometimes noisy, because the disturbance rarely outweighs the gain. A respondent with a small gain is closer to indifference, so a less favorable environment, such as lower perceived protection or more conflicting information, reduces vaccination more among inconsistent choosers in this group (implication F2). The data contain no repeated measure of perceived influenza risk, so Figure 5 and Appendix Table D4 compare groups in which the gain is plausibly smaller: Republicans, respondents with fewer than two chronic conditions, and those born after 1955. Each group differs from the others in more than perceived risk, so the comparisons show where the decline occurs, not how much of it perceived risk explains.

Partisanship. Among Republicans, receipt falls most in the bottom e^{**} tertile. Between 2014–2017 and the pooled 2021–2026 periods, adjusted receipt among Republicans falls by 15.3 percentage points (s.e. 3.1) in the bottom tertile, nearly three times the full-sample fall of 5.6 points (s.e. 1.9), and by 4.7 points (s.e. 3.2) in the top tertile, so the top-minus-bottom gap widens by 10.5 points (s.e. 4.4; Appendix Table D4; Figure 5, Panels A–B). Among Democrats, neither tertile changes significantly, and the gap changes by 4.0 points (s.e. 5.1). The bottom-tertile fall is 14.5 points larger among Republicans (s.e. 4.6), but the 6.6-point difference between the parties in gap widening (s.e. 6.7) is not statistically significant, because receipt among Republicans also falls in the top tertile. Party composition does not

Figure 5: Influenza Vaccination by e^{**} Tertile within Party, Age, and Chronic-Condition Groups



Notes: Each panel plots the covariate-adjusted change since 2014–2017 in the share reporting a flu shot, in percentage points, by fixed e^{**} tertile. Each group has its own model, which interacts e^{**} and quantitative-ability tertile indicators with reporting period and includes sex, birth year, race and ethnicity, and education with period-specific coefficients and months since the preceding interview with a single coefficient. Predictions are averaged over the group's respondents, keeping their covariates and quantitative-ability tertiles fixed across periods and setting recall length to the group mean. Tertile cutoffs are the full-sample values. Party is from UAS 71 (2016), and chronic conditions are the baseline count of diabetes, high blood pressure, cancer, heart disease, lung disease, stroke, and arthritis. N is the number of respondents in the group. Appendix Table D4 reports the pooled post-2020 changes. Bars are pointwise 95 percent confidence intervals clustered by respondent.

account for the divergence: adding party identification with pre/post-specific coefficients, which absorbs the post-2020 partisan shift in influenza vaccination (Sandlin, 2025), leaves the e^{**} change at 2.15 points per standard deviation, compared with 2.28 for the same respondents (Figure 8, Panel A). Perceived risk is only one reason the expected gain from a flu shot may have fallen among Republicans. Lower trust in public health agencies, which in the UAS predicted COVID-19 vaccine hesitancy more strongly than party identification (Choi and Fox, 2022), and spillover from COVID-19 vaccine attitudes (Lunz Trujillo et al., 2024) could lower it as well. Either could lower G in Equation (1), so under the conditions of implication F2, either could generate the pattern, and the party comparison cannot tell them apart.

Health risk. The widening is larger among respondents with fewer chronic conditions. Among them, adjusted receipt falls by 7.1 percentage points (s.e. 2.5) in the bottom tertile and rises by 4.3 points (s.e. 2.4) in the top tertile, a widening of 11.4 points (s.e. 3.5), slightly larger than the college gap (Appendix Table D4; Figure 5, Panels E–F). Among respondents with two or more conditions, neither tertile changes significantly, and the gap widens by 1.2 points (s.e. 4.3); the difference between the two groups, 10.2 points (s.e. 5.5), is significant at the 10 percent level. The age split points the same way but is smaller: the gap widens by 8.9 points (s.e. 3.8) among respondents born after 1955 and by 5.6 points (s.e. 3.7) among those born earlier (Figure 5, Panels C–D), a difference of 3.3 points (s.e. 5.3).

Less consistent respondents may also follow through less often on a decision to get vaccinated (Bronchetti et al., 2015). The experiment measures mistakes made at the moment of choice, not lapses between deciding and acting, and without repeated influenza intentions the data cannot separate the two.

5.2 COVID-19 Vaccination During the Rollout

5.2.1 Associations Across Rollout Phases and Waves

Neither measure predicts COVID-19 receipt in the initial rollout. During Waves 21–23 (December 23, 2020–February 15, 2021), 2 to 16 percent of the sample had received a dose. In the joint model, the initial association is 0.30 points per standard deviation for e^{**} and -0.55 for quantitative ability, and neither differs significantly from zero (Table 3, column 3).

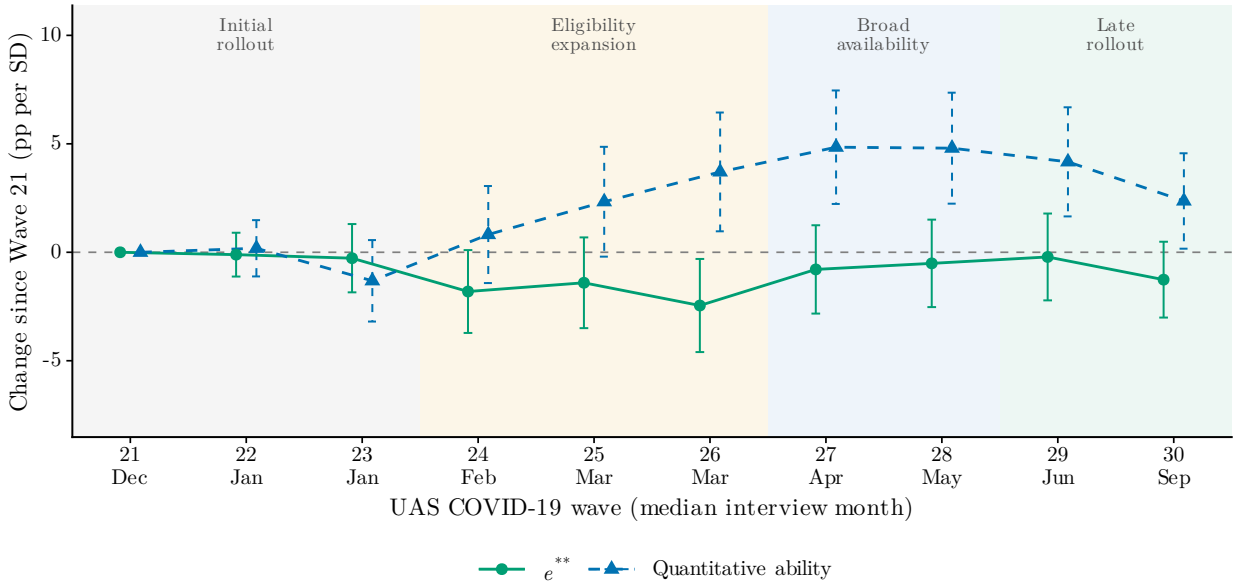
The quantitative-ability association then rises as eligibility expands. It is 2.65 points per standard deviation higher during eligibility expansion than in the initial rollout (s.e. 0.86; Table 3, column 3). The availability widening, the change from the initial rollout to broad adult availability, is 5.22 points (s.e. 1.21). This brings the association to about 4.7 points,

Table 3: Choice Consistency, Quantitative Ability, and COVID-19 Vaccine Receipt

	(1) e^{**}	(2) Quantitative ability	(3) Both
e^{**}: initial rollout	0.20		0.30
	(0.43)		(0.43)
Eligibility expansion minus initial	−1.33*		−1.75**
	(0.73)		(0.74)
Broad availability minus initial	0.32		−0.51
	(0.96)		(0.98)
Late rollout minus initial	−0.07		−0.64
	(0.88)		(0.89)
Quantitative ability: initial rollout		−0.48	−0.55
		(0.57)	(0.58)
Eligibility expansion minus initial		2.22***	2.65***
		(0.86)	(0.86)
Broad availability minus initial		5.10***	5.22***
		(1.19)	(1.21)
Late rollout minus initial		3.42***	3.57***
		(1.07)	(1.09)
Observations	19,954	19,954	19,954
Respondents	2,364	2,364	2,364
Mean outcome, initial rollout (%)	8.2	8.2	8.2
Mean outcome, eligibility expansion (%)	43.7	43.7	43.7
Mean outcome, broad availability (%)	75.3	75.3	75.3
Mean outcome, late rollout (%)	81.6	81.6	81.6
Adjusted R^2	0.447	0.448	0.449

Notes: Entries are associations with cumulative self-reported receipt of at least one COVID-19 dose, in percentage points per standard deviation of each measure, from unweighted linear probability models of Equation (8). For each measure, the first row is the association in the initial rollout and the other rows are changes from the initial rollout. The phases are the initial rollout (Waves 21–23), eligibility expansion (24–26), broad adult availability (27–28), and late rollout (29–30). Columns 1 and 2 enter one measure at a time; column 3 enters both. All models include wave fixed effects, phase-specific coefficients on sex, birth year, race and ethnicity, six education categories, and baseline state, and wave-specific slopes in interview date. e^{**} is scaled by its standard deviation across unique respondents in the estimation sample, and quantitative ability by its standard deviation in the construction sample. The mean outcome is the percentage of respondent-wave observations in the phase with at least one dose. Appendix Table E1 reports wave-specific estimates and Appendix Table F1 the estimates with e^{*}. Standard errors clustered by respondent are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Figure 6: COVID-19 Vaccine Receipt by Survey Wave



Notes: The figure plots the change since Wave 21 in the associations of e^{**} and quantitative ability with cumulative self-reported receipt of at least one COVID-19 dose, in percentage points per standard deviation, from the joint model in Equation (9). The model includes wave fixed effects, wave-specific coefficients on sex, birth year, race and ethnicity, education, and baseline state, and wave-specific slopes in interview date. Both measures are scaled as in Table 3. The sample has 19,954 respondent-wave observations from 2,364 respondents. Wave 21 is the reference and has no confidence interval. Shading marks the rollout phases defined in Table 3. Appendix Table E1 reports the estimates. Bars are pointwise 95 percent confidence intervals clustered by respondent.

roughly a quarter of the 20-point college gap in receipt in that phase (Appendix Table D2, Panel B). By wave, the change relative to Wave 21 stays near zero through Wave 24 and reaches about 4.8 points in Waves 27 and 28. Wave 27 spans April 19, the date by which all adults were eligible (Diesel et al., 2021). By Wave 30, the change falls to 2.4 points, significantly below its Wave 27 value ($p = 0.03$; Figure 6; Appendix Table E1), consistent with catch-up in a cumulative outcome once most respondents had accepted vaccination and could obtain a dose (implication C3).

The e^{**} association moves the other way during eligibility expansion. It falls by 1.75 points per standard deviation (s.e. 0.74; Table 3, column 3). In later phases, the change from the initial rollout is small and not significant (Figure 6). One possible explanation follows the logic of implication F1. More consistent respondents make fewer mistakes relative to their own assessment of benefits and risks, so if waiting looked like the better choice for their health during this phase, they would delay more than less consistent respondents, and the return of the association would mean they caught up once vaccination looked worthwhile again. The ten-day federal pause of the Janssen vaccine in April 2021 (MacNeil et al., 2021) is one event that could have made waiting look better. The timing fits only in part: the

decline already appears in Wave 24 (−1.81, s.e. 0.97), two months before the pause, and is small in Wave 27, which spans it (−0.79, s.e. 1.04; Appendix Table E1). The data cannot test this explanation.

5.2.2 What Accounts for the Availability Widening?

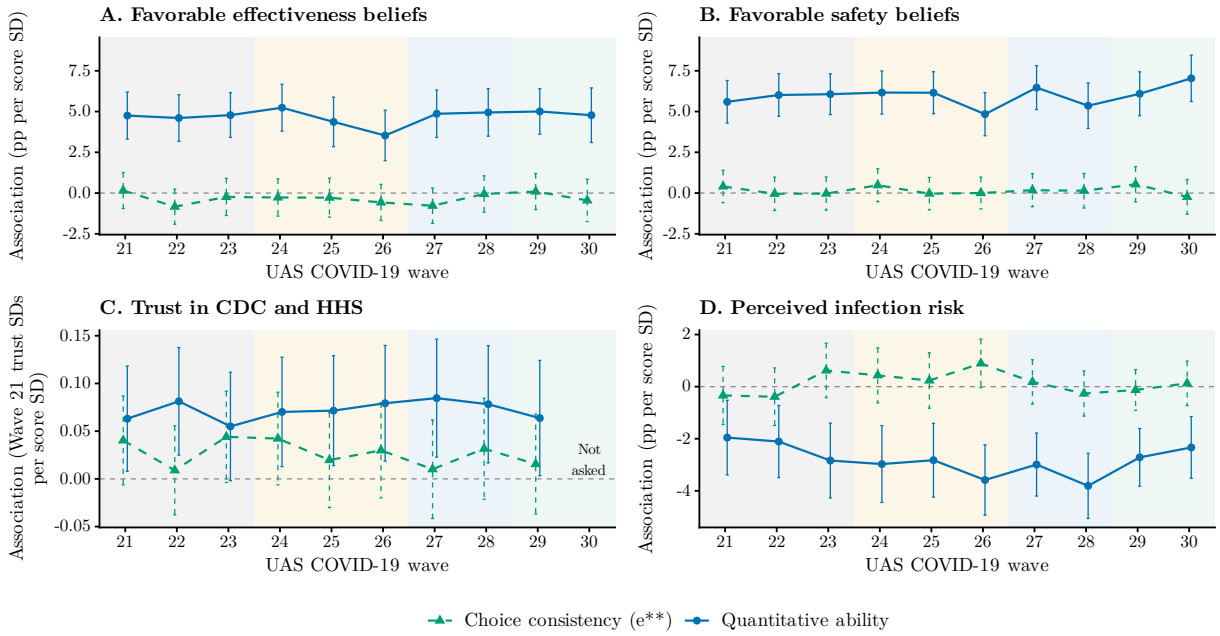
Favorable beliefs about the vaccines that higher-ability respondents already held in Wave 21 account for about 60 percent of the availability widening. After adjusting for those beliefs, lower-ability respondents also more often reported not knowing how to book or where to go, by 2.8 percentage points ($p = 0.064$).

Three explanations could produce the availability widening, and all three imply that the gradient opens only once access opens, because receipt waits for the later of acceptance and access (Equation (4)). If higher-ability respondents learned faster about the vaccines during the rollout (implication C1), the ability gap in favorable beliefs should widen before or as the receipt gap opens. If they held more favorable views from the start, the gap should already be present in Wave 21, and adding those views should shrink the widening. If they acted sooner on the same views or arranged a dose more easily once eligible (implication C2), part of the widening should remain after the views are held fixed, and lower-ability respondents with the same views should more often report trouble getting a dose.

Views during the rollout. Beliefs about the vaccines and trust in health agencies changed little while most adults became eligible; perceived infection risk fell (Appendix Figure E1). Between Waves 21 and 29, mean favorable effectiveness beliefs rose from about 74 to about 79 on the 0–100 scale, and favorable safety beliefs from about 79 to about 83. The CDC/HHS trust index stayed near 2.4 on the 1–4 scale. Perceived three-month infection risk fell from about 23 to about 10 percent. In Wave 30 (September–October 2021), effectiveness beliefs fell to about 64 and infection risk rose to about 16. These are means among respondents observed in each wave, not within-person changes.

Who held which views. Higher-ability respondents viewed the vaccines more favorably, trusted the CDC and HHS more, and perceived less risk of infection from the first wave (Figure 7). In Wave 21, when 1.9 percent of the sample had a dose, favorable effectiveness and safety beliefs are 4.8 and 5.6 points higher per standard deviation of quantitative ability, the trust index is 0.063 of its Wave 21 standard deviation higher (s.e. 0.028), and perceived infection risk is about 2.0 points lower (s.e. 0.7). The vaccine-belief gradients do not widen as access expands. By Wave 27, the safety-belief gradient has changed by 0.9 points (s.e. 0.7), which rules out increases larger than 2.2 points, the effectiveness gradient by 0.1 points,

Figure 7: COVID-19 Vaccine Views by Survey Wave: Associations with e^{**} and Quantitative Ability



Notes: Each panel plots, wave by wave, the association of e^{**} and of quantitative ability with one view of COVID-19 or its vaccines. Points are levels in each wave, not changes from Wave 21. Panels A and B are favorable effectiveness and safety beliefs, 100 minus the stated probability that a vaccinated person will contract COVID-19 or that vaccination will cause serious side effects. Panel D is the stated probability (0–100) of infection in the next three months, so higher values mean more perceived risk. In these three panels, associations are in points per standard deviation of each measure. Panel C is the CDC/HHS trust index, the mean of two items each scored 1–4 and requiring both, standardized by its Wave 21 mean and standard deviation; its associations are in Wave 21 standard deviations of trust per standard deviation of each measure. Trust was not asked in Wave 30. Panels A–D use 20,167, 20,157, 18,032, and 20,182 respondent-wave observations from 2,377, 2,377, 2,305, and 2,378 respondents. Each panel comes from one model with wave fixed effects and wave-specific coefficients on both measures, sex, birth year, race and ethnicity, education, baseline state, and interview date. Quantitative ability is in construction-sample standard deviations, and e^{**} is scaled by its standard deviation across unique respondents in each panel's sample. Shading marks the rollout phases. Bars are pointwise 95 percent confidence intervals based on standard errors clustered by respondent. Appendix Table E2 reports the estimates.

and the trust gradient by 0.022 standard deviations (s.e. 0.025; Appendix Table E2). Faster learning during the rollout would have widened the belief gap before or with the receipt gap (implication C1), and that does not happen. The e^{**} associations are close to zero for all four views in every wave. These wave-specific gradients do not trace individual revisions or show when the initial differences formed.

Accounting for the widening. Wave 21 vaccine beliefs account for most of the widening, trust for less, and perceived infection risk for none (Table 4). On the common views sample, the availability widening is 5.27 points (s.e. 1.24). Adding the two beliefs with phase-specific

Table 4: Changes in the Quantitative-Ability Association with COVID-19 Receipt after Adding Wave 21 Views

	(1) Expansion minus initial	(2) Broad availability minus initial	(3) Respondents
Full sample	2.65***	5.22***	2,364
	(0.86)	(1.21)	
Views sample, no added views	2.52***	5.27***	2,049
	(0.90)	(1.24)	
Views sample + trust	1.79**	4.03***	2,049
	(0.88)	(1.18)	
Views sample + beliefs	0.49	1.99	2,049
	(0.89)	(1.25)	
Views sample + infection risk	2.70***	5.69***	2,049
	(0.89)	(1.23)	

Notes: Each entry is the change, relative to the initial rollout (Waves 21–23), in the association between quantitative ability and cumulative receipt of at least one COVID-19 dose, in percentage points per standard deviation of quantitative ability. Column 1 is the change during eligibility expansion (Waves 24–26), column 2 the change during broad adult availability (Waves 27–28), which is the availability widening, and column 3 the number of respondents. The first row repeats Table 3, column 3. The other rows use the 2,049 respondents who answered both trust items and both vaccine beliefs in Wave 21 (18,876 respondent-wave observations); all of them also report perceived infection risk, and 39 had a dose by Wave 21. Each of the last three rows adds one set of Wave 21 views, interacted with rollout phase: the CDC/HHS trust index, the mean of two items each scored 1–4; favorable effectiveness and safety beliefs, 100 minus the stated probability that a vaccinated person will contract COVID-19 or that vaccination will cause serious side effects; or the stated probability (0–100) of infection in the next three months. Each view is standardized by its mean and standard deviation among all Wave 21 respondents who answered it. All rows estimate Equation (8) with e^{**} and quantitative ability entered jointly and the controls of Table 3. Standard errors clustered by respondent are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

coefficients lowers it to 1.99 (s.e. 1.25), a reduction of about 60 percent, and lowers the eligibility-expansion change from 2.52 to 0.49 points. Neither remaining change is statistically significant. The confidence interval for the availability widening reaches about 4.4 points, so a sizable remainder cannot be ruled out. Adding institutional trust instead lowers the widening to 4.03 points (s.e. 1.18), about a quarter less. Adding perceived infection risk raises it to 5.69 points (s.e. 1.23): higher-ability respondents perceived less risk of infection yet vaccinated earlier. The views were measured before most respondents could get a dose, but these comparisons are accounting exercises, not estimates of a mediating channel; the views may capture other traits correlated with ability.

Reported barriers. After adjusting for Wave 21 beliefs, lower-ability respondents more often reported not knowing how to book or where to go, but no more obstacles after trying

(Table 5). The remaining widening could reflect acting sooner or arranging a dose more easily (implication C2); questions in Waves 25–29, whose median interview months run from March to June 2021, provide a partial check. Respondents not yet vaccinated were asked whether they had tried to get a dose; those who had tried scheduling were asked about obstacles, and those who did not report trying were asked why (items `cr084`–`cr088`; Appendix B.7). I fix the cohort at the 2,013 respondents unvaccinated in Wave 21 with both baseline beliefs observed and record whether each reason was reported at least once before the first reported dose. The 490 vaccinated by Wave 24 were never asked these items and count as not reporting them.

With baseline beliefs included, the bottom-minus-top quantitative-ability difference in not knowing how to book or where to go is 2.8 points (s.e. 1.5; $p = 0.064$); the adjusted share in the bottom tertile is 5.9 percent, nearly twice the 3.1 percent in the top. The difference in belief reasons moves from 4.1 points (s.e. 2.5) without belief controls to −4.0 points (s.e. 2.5) with them; neither difference is statistically significant. Obstacles after trying are reported by about 1–2 percent in each tertile. Lower-ability respondents remained unvaccinated longer and had more chances to answer the reason questions, 1.97 against 1.26 on average. The 2.8-point gap may therefore overstate a difference in arranging a dose and does not identify one.

Most respondents who reported only not knowing how to book or where to go subsequently reported vaccination, whereas most reporting only belief reasons remained unvaccinated by Wave 29 (Table 6). Of the 60 who gave only the former reason, 39 (65 percent) were vaccinated by Wave 29; of the 290 who gave only belief reasons, 43 (15 percent) were vaccinated. Of the 26 who reported only an obstacle after trying, 23 were vaccinated. These counts refer to unique respondents, not repeated answers, and do not establish that the reported reason caused a delay.

6 Measurement Error and Robustness

The availability widening holds up in every check, and the flu divergence is less robust. Figure 8 reports both estimates across specifications. The availability widening stays well above zero in every row. The flu divergence ranges from 1.60 to 2.84 points in all but one check, but its 95 percent confidence interval includes zero in 8 of the 18 checks, and it falls to 0.93 points when fast respondents are excluded at the stricter cutoff (Section 6.2). For rows that add a control or an instrument, open points show the corresponding OLS baseline on the same observations, so the gap between the two points reflects the added control or IV correction rather than a change in sample. Rows that restrict the sample have no such

Table 5: Reported Barriers to COVID-19 Vaccination by Quantitative-Ability Tertile

	Baseline controls			+ Wave 21 beliefs		
Reported in Waves 25–29	(1) Bottom (%)	(2) Top (%)	(3) Bottom minus top	(4) Bottom (%)	(5) Top (%)	(6) Bottom minus top
Belief reason	18.5	14.4	4.06	14.3	18.3	-4.00
			(2.53)			(2.47)
Did not know how to book or where to go	6.3	2.8	3.52**	5.9	3.1	2.80*
			(1.47)			(1.51)
Obstacle after trying to book	1.1	1.5	-0.45	1.1	1.5	-0.37
			(0.84)			(0.88)
Respondents	2,013			2,013		

Notes: Columns 1, 2, 4, and 5 are the adjusted percentages of respondents in the bottom and top quantitative-ability tertiles who reported each reason at least once in Waves 25–29 before their first reported dose. Columns 3 and 6 are the bottom-minus-top differences, in percentage points. Adjusted percentages are predictive margins: the mean predicted probability when every respondent is assigned to the given tertile and other covariates keep their observed values. The sample is the 2,013 respondents who were unvaccinated in Wave 21 and reported both Wave 21 vaccine beliefs. It includes 490 respondents vaccinated by Wave 24, who were never asked the reason questions and are coded as reporting no reason. Belief reasons are that the vaccine is not effective, not safe, or not necessary. Not knowing how to book or where to go was asked only of respondents who did not report trying to get a dose. Obstacles after trying are a hard-to-understand appointment system, technical problems, or a missed appointment (Appendix B.7). Reasons can overlap, and a zero means no recorded report. On average, bottom-tertile respondents answered the reason questions 1.97 times and top-tertile respondents 1.26 times. Each row comes from a person-level linear probability model with tertile indicators (middle tertile omitted), continuous e^{**}, sex, birth year, race and ethnicity, six education categories, baseline state, and Wave 21 interview timing; columns 4–6 add continuous Wave 21 safety and effectiveness beliefs. Tertile cutoffs are those of Appendix Table D3. Heteroskedasticity-robust standard errors are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table 6: Reported Reasons and COVID-19 Vaccination Status by Wave 29

Reported reason	(1) People	(2) Vaccinated by Wave 29	(3) Unvaccinated in Wave 29	(4) Status unresolved
Panel A. Each reason counted once per person; reasons may overlap				
Did not know how to book or where to go	87	44	36	7
Vaccine ineffective, unsafe, or unnecessary	316	48	243	25
Obstacle after a booking attempt	27	23	3	1
Panel B. Mutually exclusive histories among respondents reporting a reason				
Only did not know how to book or where to go	60	39	15	6
Only vaccine ineffective, unsafe, or unnecessary	290	43	222	25
Only obstacle after a booking attempt	26	23	3	0
Both how/where-to-book and belief reasons	26	5	21	0
Both how/where-to-book and post-attempt obstacles	1	0	0	1
Any of the three reasons	403	110	261	32

Notes: Entries are numbers of respondents, without regression adjustment. Column 1 counts the respondents who reported each reason, and columns 2–4 divide them by status in Wave 29: vaccinated means a reported dose by Wave 29, unvaccinated means a negative report in Wave 29 and no earlier positive report, and unresolved covers everyone else. Panel A counts each of the three reasons in Table 5 separately, so its rows overlap and do not sum. Panel B assigns each respondent to one combination of the three reasons; "only" refers to these three reasons, not to other survey reasons, and no respondent reported both a belief reason and an obstacle after a booking attempt. The sample is the Table 5 cohort of 2,013 respondents; the 1,610 who reported none of the three reasons, including those vaccinated before Wave 25, are not shown. A reason is recorded if it was reported at least once in Waves 25–29 before the first reported dose, and repeated reports count once. Appendix B.7 defines the reasons.

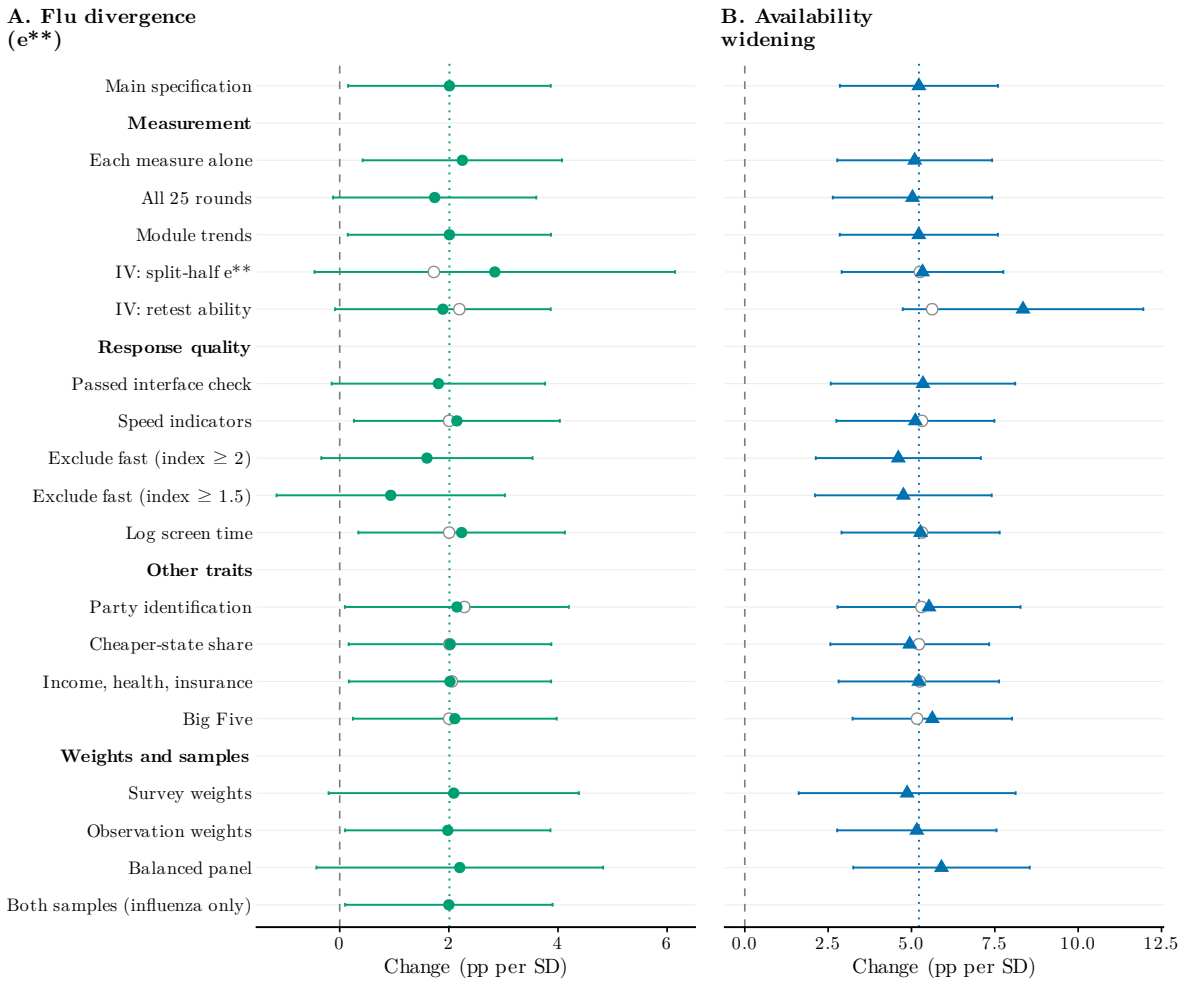
baseline. Appendix Table F4 reports every estimate.

6.1 Measurement Error

Classical error. Both measures contain error, and the choice experiment was administered only once. Error is classical if it is unrelated to true ability and to other determinants of vaccination. Such error places people with different true abilities at the same measured score. In a single-measure model, it shrinks the measure's associations, and their changes over time, toward zero by a factor related to its reliability. Three of the four ability components have test-retest correlations of 0.74 to 0.77 between UAS–HRS Waves 1 and 2, and the composite has a correlation of 0.82 among the 2,344 respondents with all four components in both waves (Appendix Table B4). For the consistency indices, only split-half correlations are available, 0.59 for e^{**} and 0.41 for e^{*}. The Spearman–Brown formula puts the reliability of the full 25-round indices at 0.75 and 0.58 (Appendix Table B4). These figures come from a single sitting and leave out changes from one occasion to the next, so reliability across years may be lower. Because consistency is measured less reliably, the shrinkage works against the flu divergence, and it could hide a consistency association with COVID-19 receipt. Such attenuation cannot by itself produce the opposite-signed changes during eligibility expansion, because shrinking toward zero does not reverse a sign. These conclusions assume that the sample and each measure's relation to perceived gains are stable over time. An unbalanced panel of aging respondents need not meet these conditions, and the direction of any resulting bias is not established.

Error that crosses measures. In joint models, error in one measure leaves part of the true ability it captures uncontrolled. Because the two measures are correlated, that part can load onto the other measure and push its coefficient toward its single-measure value. Comparing single-measure and joint models shows how sensitive the estimates are to conditioning on the other measured score. The single-measure models separate the two measures for COVID-19 but not for influenza (Tables 2 and 3, columns 1–2). With quantitative ability alone, the availability widening is 5.10 points (s.e. 1.19), against 5.22 in the joint model; with e^{**} alone, the corresponding e^{**} change is 0.32 points (s.e. 0.96). With e^{**} alone, the flu divergence is 2.25 points (s.e. 0.93), and with quantitative ability alone, its flu change is 2.26 points (s.e. 1.20). The linear flu changes are about the same size for the two measures, alone or together. The difference between them shows up in the tertile paths, where receipt falls mainly in the bottom e^{**} tertile and by similar amounts across quantitative-ability tertiles (Section 5.1.2). The consistency gap also widens in unadjusted pre/post rates that do not control for quantitative ability (Appendix Figure D1).

Figure 8: Sensitivity of the Principal Changes



Notes: The figure reports the two principal estimates under alternative specifications, in percentage points per standard deviation. Panel A is the flu divergence, the change in the e^{**} association with reported influenza vaccination from 2014–2019 to 2021–2026 (Equation (6)). Panel B is the availability widening, the change in the quantitative-ability association with COVID-19 receipt from the initial rollout to broad adult availability (Equation (8)). Filled points are the estimates; open gray points are OLS estimates on the same observations, shown for rows that add a control or an instrument. Measurement rows enter each measure alone; keep respondents with all 25 usable rounds; add module-specific trends; instrument each odd- or even-round half of e^{**} with the other half and stack both directions; or instrument the Wave 1 quantitative-ability composite with its UAS–HRS Wave 2 retest, whose OLS baseline also uses Wave 1 ability. Response-quality rows keep respondents who passed the final paid interface task; add indicators for fast and slow responses; drop respondents with a speed index (Section 2.4) of at least 2 or at least 1.5; or add log mean screen time. The speed- and screen-time rows omit respondents without the required timing data. Other-trait rows add party identification (UAS 71); the mean budget share allocated to the cheaper state; household income, self-rated health, chronic conditions, and insurance; or the Big Five scores. Weight and sample rows use UAS survey weights, inverse-probability-of-observation weights, a balanced panel, or, for influenza only, the respondents who are also in the COVID-19 sample. Except in the single-measure row, every model enters e^{**} and quantitative ability jointly, and added controls have pre/post-specific coefficients for influenza and phase-specific coefficients for COVID-19. Bars are 95 percent confidence intervals based on standard errors clustered by respondent. The dashed line marks zero, and the dotted line marks the main estimate. Appendix Table F4 reports the estimates, same-sample baselines, and sample sizes.

Instrumental-variable corrections. I compute e^{**} on odd and even rounds and stack both instrumenting directions, instrumenting each half with the other and clustering by respondent, as in the ORIV approach of Gillen et al. (2019). On the same observations, the flu divergence rises from 1.73 points in OLS to 2.84 points (s.e. 1.69), and its confidence interval includes zero; the availability widening is 5.33 points (s.e. 1.24), against 5.25. Instrumenting the Wave 1 ability composite with its pre-pandemic Wave 2 retest raises the availability widening from 5.62 to 8.35 points (s.e. 1.84) on the same observations and leaves the flu divergence at 1.89 (Figure 8). Each correction raises the coefficient on the instrumented measure, as attenuation predicts. Both require independent errors across the two measurements, and real changes between waves or common response tendencies would violate that.

Ability components. No single test drives the availability widening. Entered one at a time with e^{**}, each of the four components has a positive and significant broad-availability change, from 2.83 to 3.52 points (Appendix Table F2, Panel B). The influenza changes are less uniform, from -0.18 points for financial literacy to 2.17 for numeracy, and the flu divergence stays between 2.00 and 2.32 points whichever component is entered (Panel A). The components are correlated, so these separate models do not identify independent channels.

6.2 Response Quality and Index Construction

Response quality. Inattentive responses could produce low consistency scores that also correlate with vaccination. The flu divergence depends on how fast respondents are handled. Excluding respondents with a speed index of at least 2, the fast group defined in Section 2.4, gives 1.60 points (s.e. 0.99). A stricter cutoff of 1.5, roughly two-thirds of the median screen time for uniform speeding across screens, removes 21.3 percent of influenza respondents and gives 0.93 points (s.e. 1.07). This lower cutoff tests sensitivity to moderately fast responses as well as extreme speeding; it is not a validated threshold for inattention. Keeping only respondents who passed the final paid interface task gives 1.81 points (Figure 8, Panel A). Adding fast- and slow-response indicators with period-specific coefficients instead leaves the estimate at 2.15 points. So speed itself does not account for the divergence, but much of the estimate depends on respondents who answered quickly. Fast respondents have a mean e^{**} of 0.60, against 0.76 among respondents with typical speed (Section 2.4), and the divergence comes from the bottom of the e^{**} distribution, so excluding them removes part of the group that drives the result. If consistency reflects attention to each choice, fast answering may be part of what it measures. These checks cannot tell whether the low scores of fast respondents reflect a lasting tendency to choose with little attention or carelessness confined to this survey. The availability widening changes little and is 4.76 points (s.e. 1.35) at the stricter cutoff.

Index construction. Adding module-specific time trends leaves the flu divergence at 2.01 points. Keeping only respondents with all 25 usable rounds, which avoids the tendency of incomplete records to raise the CCEI (Choi et al., 2014), lowers it to 1.74 points (s.e. 0.95), and its confidence interval includes zero. Of the 3,134 experiment respondents, 3,066 have complete records; in the joint influenza sample, this restriction retains 2,874 of 2,939 respondents. The change comes from a smaller point estimate rather than lost precision. The availability widening is 5.22 points with module trends and 5.03 points with complete rounds. With e^* in place of e^{**}, the flu change is 1.62 points (s.e. 1.00) conditional on quantitative ability and 1.83 points (s.e. 0.98) without it, and the availability widening is 5.19 points (s.e. 1.21; Appendix Table F1).

6.3 Other Traits

The measures may proxy for preferences or circumstances that also predict vaccination. Adding the mean budget share allocated to the cheaper state, a proxy for risk attitudes, leaves the flu divergence at 2.02 points and the availability widening at 4.95 (Figure 8). Adding the Big Five scores gives 2.11 and 5.62 points. Adding income, self-rated health, chronic conditions, and insurance leaves both close to their main values. Section 4.3 leaves these four out of the main models because they may partly reflect the two measures, so conditioning on them could absorb part of the associations; the estimates barely move. Party identification, discussed in Section 5.1.3, leaves the flu divergence at 2.15 points, against 2.28 on the same respondents. These controls are not exhaustive measures of preferences or circumstances.

6.4 Sample Composition and Weights

Attrition. Selective observation is a concern in the unbalanced influenza panel. Higher quantitative ability predicts being observed: a one-standard-deviation increase lowers the probability of having no interview in a reporting period by 1.57 percentage points (s.e. 0.57), against a mean of 25.9 percent. Neither consistency index predicts absence, and item nonresponse to the flu-shot question is 0.1 percent (Appendix Table F3). Attrition is therefore more of a concern for the quantitative-ability estimates than for the consistency estimates.

Weights. The inverse-probability-of-observation weights correct for selection only if vaccination and observation are independent given the weighting-model predictors. With these weights, the flu divergence is 1.98 points and its confidence interval excludes zero. With UAS survey weights, it is 2.09 points but less precise, and the interval includes zero. The

availability widening is 5.16 points with observation weights and 4.87 points with survey weights (Figure 8).

Balanced panels and common samples. The balanced influenza panel ends in 2023–2025, so it cannot evaluate the larger final-period estimate. In that panel, the flu divergence is 2.20 points, with a confidence interval that includes zero. Restricting the influenza analysis to respondents in the COVID-19 sample gives 2.00 points. These comparisons and the weights address observed composition, not selection on unobserved vaccination propensities.

7 Conclusion

7.1 Summary of Findings

Test-based quantitative ability and revealed-preference choice consistency are not interchangeable measures of decision-making ability. Quantitative ability predicts who got a new vaccine earlier, and choice consistency predicts whose receipt of a familiar vaccine fell.

For influenza, neither measure predicts vaccination on average, but the choice-consistency association rises over time and is larger after 2020, with some separation already in 2017–2019. Between 2014–2017 and 2025–2026, adjusted receipt falls by 9.4 percentage points in the bottom consistency tertile and changes little in the top. In both the bottom and top quantitative-ability tertiles, it falls by about 4 points. Receipt in the bottom consistency tertile falls more among Republicans than among Democrats, and the gap grows more among respondents with fewer chronic conditions. Adding party identification leaves the linear divergence almost unchanged.

During the COVID-19 rollout, neither measure predicts receipt at first. Quantitative ability then becomes more positively associated with receipt as eligibility expands, reaching about 4.7 percentage points per standard deviation during broad adult availability. Choice consistency instead becomes temporarily less positively associated during eligibility expansion. Higher-ability respondents already held more favorable safety and effectiveness beliefs in the first wave, and these belief gradients did not widen. Adding these beliefs reduces the later widening in receipt by about 60 percent and leaves neither phase change significant; adding institutional trust reduces it by about a quarter, and adding perceived infection risk does not reduce it. After adjusting for baseline beliefs, lower-ability respondents more often reported not knowing how to book or where to go, and most of those who gave only this reason were vaccinated by Wave 29. These patterns are suggestive evidence that higher-ability respondents acted on more favorable prior assessments once they could get a dose, with a

smaller and less precise role for knowing how to get one.

The results also matter for how these measures are validated. Both measures have small average associations with influenza vaccination, yet consistency predicts where receipt falls over time. A measure checked only against the level of one outcome can therefore look uninformative and still carry information about change. Other work finds that experimental consistency corresponds only weakly to consistency in supermarket purchases (Chen et al., 2025), and that consistency measures and cognitive tests can rank populations differently (Cappelen et al., 2023).

7.2 Limitations

The limitations mainly come from the data. First, the consistency indices come from a single administration of the choice experiment, fielded before most outcomes were measured, so their test-retest reliability is unknown. The split-half analysis in Section 2.4 measures reliability within one sitting but does not establish stability across years. The flu divergence also falls by more than half when fast respondents are excluded at the stricter cutoff (Section 6.2), so the indices may capture carelessness specific to this survey as well as a lasting tendency to make mistakes in repeated decisions. Second, both vaccination outcomes are self-reported, and neither is precisely dated. Each influenza report covers about two seasons, so individual seasons cannot be separated, and COVID-19 receipt is dated by interview rather than by dose, so recorded receipt may lag actual vaccination. Third, the sample is older than the U.S. adult population, and 76.5 percent of respondents are non-Hispanic White, so the results describe adults for whom influenza vaccination is common. Last but not least, the data cannot fully separate the mechanisms behind the influenza and COVID-19 patterns. For influenza, the linear changes are about the same size for the two measures, so the distinction between them rests on where in each distribution receipt falls (Sections 5.1.2 and 6.1). Influenza beliefs and intentions are not observed in each period, so lapses between deciding and acting could produce the pattern that the framework attributes to mistakes at the moment of choice. For COVID-19, Wave 21 beliefs may also capture other traits correlated with ability, including skills that make it easier to book a dose. The barrier questions were asked only of respondents not yet vaccinated, so lower-ability respondents, who stayed unvaccinated longer, had more chances to report a barrier, and the reported-barrier comparison does not identify a difference in arranging a dose. Error that is not classical and changes in sample composition can also affect the interaction estimates, so the opposite-signed changes during the rollout do not by themselves establish distinct latent mechanisms.

7.3 Policy Implications

The two vaccination problems may call for different tools, but the UAS estimates do not evaluate interventions.

For a familiar vaccine, the adults drifting away are hard to find with standard targeting. Receipt fell most among less consistent Republicans, and the consistency gap grew most among respondents with fewer chronic conditions; the expected gain from a flu shot may be small in both groups. Consistency is only weakly related to age, sex, education, and income (Appendix Table B3) and to quantitative ability (correlation 0.27), so outreach targeted on these characteristics may miss many of these adults. Appeals to protecting others and to social norms do not rely on perceived personal risk, so they may still reach them. Prosociality measured in incentivized experiments predicts protective health behaviors such as physical distancing during the pandemic (Campos-Mercade et al., 2021), and altruistic motives enter influenza vaccination decisions (Shim et al., 2012). In experiments, messages describing people harmed when others do not vaccinate raised influenza vaccination intentions, especially among those not vaccinated in the previous year (Li et al., 2016). Prosocial and social-image messages raised COVID-19 vaccination intentions across subgroups (James et al., 2021), and accurate information on how many others intended to be vaccinated raised intentions, particularly among those initially unsure (Moehring et al., 2023). Among strongly hesitant adults, however, personal-benefit messages reduced COVID-19 hesitancy more than collective-benefit messages (Freeman et al., 2021).

Who delivers these messages may matter as much as what they say, because the groups drifting away include adults who trust federal health agencies less. Partisans are less likely to view vaccines as safe and to vaccinate their children when the other party holds the presidency, and mediation estimates link this shift to lower trust in government (Krupenkin, 2021). Trust in the CDC as a source of vaccine information fell from 63 percent of adults in September 2023 to 47 percent in January 2026 (KFF, 2026). Messengers that a group already trusts can move it. Among unvaccinated Republicans, endorsements by Republican leaders raised COVID-19 vaccination intentions by 7.0 percent relative to endorsements by Democratic leaders and by 5.7 percent relative to a neutral control, while Democratic endorsements made Republicans less willing to encourage others to vaccinate (Pink et al., 2021). In a randomized YouTube campaign with 1,014 treated counties, an advertisement featuring Donald Trump's endorsement of the vaccine on Fox News raised COVID-19 vaccinations by about 103 doses per treated county, about 104,000 in total, although the estimate is significant only at the 10 percent level in a one-tailed test (Larsen et al., 2023). Physicians and nurses are another trusted source: their video messages reduced holiday travel and COVID-19 infections during the 2020 holiday season (Breza et al., 2021), but similar

messages did not change vaccination during the Omicron wave (Ho et al., 2023). So it is worth testing whether prosocial and normative messages delivered by clinicians and by leaders these adults trust suit adults drifting away from a routine vaccine better than adults who reject vaccination outright. Most of this evidence concerns stated intentions or COVID-19, so whether such messages, or universal prompts such as planning prompts (Milkman et al., 2011), raise influenza receipt among the least consistent adults is a question for field experiments.

For a new vaccine, the evidence points mainly to beliefs. Beliefs held before most adults could get a dose account for about 60 percent of the ability gap in COVID-19 receipt, so lowering the cost of acting alone may not close it, and clear, early information on safety and effectiveness may matter more. Pre-booked appointments and reminders can still help with the rest, including the lower-ability adults who reported not knowing how to book or where to go: pre-booked appointments narrow cognitive-ability differences (Elinder et al., 2023), and reminders increase receipt (Dai et al., 2021; Milkman et al., 2024).

Two recent changes in federal recommendations provide tests in other settings. RSV vaccines became available to adults aged 60 and older in 2023 under shared clinical decision-making (Melgar et al., 2023), and in 2024 a simpler age- and risk-based rule replaced that recommendation (Britton et al., 2024). If quantitative ability captures how readily people evaluate a new option, its association with RSV uptake should emerge in 2023 and narrow after the 2024 simplification. This would be a test outside a pandemic and without phased eligibility. COVID-19 recommendations moved in the opposite direction, from routine vaccination for all adults (Panagiotakopoulos et al., 2024) to individual-based decision-making in October 2025 (Centers for Disease Control and Prevention, 2025). If consistency captures how often people make mistakes in repeated decisions, it should predict which regular recipients stop getting vaccinated, and uptake should decline most at the bottom of the e^{**} distribution, particularly among Republicans and adults at lower health risk.

References

Abaluck, J. and Gruber, J. (2011). Choice inconsistencies among the elderly: Evidence from plan choice in the Medicare Part D program. *American Economic Review*, 101(4):1180–1210.

Afriat, S. N. (1967). The construction of utility functions from expenditure data. *International Economic Review*, 8(1):67–77.

Afriat, S. N. (1973). On a system of inequalities in demand analysis: An extension of the classical method. *International Economic Review*, 14(2):460–472.

Alattar, L., Messel, M., and Rogofsky, D. (2018). An introduction to the Understanding America Study internet panel. *Social Security Bulletin*, 78(2):13–28.

Allcott, H., Boxell, L., Conway, J., Gentzkow, M., Thaler, M., and Yang, D. Y. (2020). Polarization and public health: Partisan differences in social distancing during the coronavirus pandemic. *Journal of Public Economics*, 191:104254.

Alsan, M. and Eichmeyer, S. (2024). Experimental evidence on the effectiveness of nonexperts for improving vaccine demand. *American Economic Journal: Economic Policy*, 16(1):394–414.

Angrisani, M. and Kapteyn, A. (2024). Financial literacy in the Understanding America Study. *Journal of Financial Literacy and Wellbeing*, 2(1):1–20.

Apesteguia, J. and Ballester, M. A. (2015). A measure of rationality and welfare. *Journal of Political Economy*, 123(6):1278–1310.

Banerjee, A., Chandrasekhar, A. G., Dalpath, S., Duflo, E., Floretta, J., Jackson, M. O., Kannan, H., Loza, F., Sankar, A., Schrimpf, A., and Shrestha, M. (2025). Selecting the most effective nudge: Evidence from a large-scale experiment on immunization. *Econometrica*, 93(4):1183–1223.

Batty, G. D., Deary, I. J., Fawns-Ritchie, C., Gale, C. R., and Altschul, D. (2021). Pre-pandemic cognitive function and COVID-19 vaccine hesitancy: Cohort study. *Brain, Behavior, and Immunity*, 96:100–105.

Bhargava, S., Loewenstein, G., and Sydnor, J. (2017). Choose to lose: Health plan choices from a menu with dominated options. *The Quarterly Journal of Economics*, 132(3):1319–1372.

Black, C. L., Kriss, J. L., Razzaghi, H., et al. (2023). Influenza, updated COVID-19, and respiratory syncytial virus vaccination coverage among adults—united states, fall 2023. *MMWR Morbidity and Mortality Weekly Report*, 72(51):1377–1382.

Boulier, B. L., Datta, T. S., and Goldfarb, R. S. (2007). Vaccination externalities. *The B.E. Journal of Economic Analysis & Policy*, 7(1):1–27.

Brewer, N. T., Chapman, G. B., Gibbons, F. X., Gerrard, M., McCaul, K. D., and Weinstein, N. D. (2007). Meta-analysis of the relationship between risk perception and health behavior: The example of vaccination. *Health Psychology*, 26(2):136–145.

Breza, E., Stanford, F. C., Alsan, M., Alsan, B., Banerjee, A., Chandrasekhar, A. G., Eichmeyer, S., Glushko, T., Goldsmith-Pinkham, P., Holland, K., Hoppe, E., Karnani, M., Liegl, S., Loisel, T., Ogbu-Nwobodo, L., Olken, B. A., Torres, C., Vautrey, P.-L., Warner, E. T., Wootton, S., and Duflo, E. (2021). Effects of a large-scale social media advertising campaign on holiday travel and COVID-19 infections: A cluster randomized controlled trial. *Nature Medicine*, 27(9):1622–1628.

Britton, A., Roper, L. E., Kotton, C. N., Hutton, D. W., Fleming-Dutra, K. E., Godfrey, M., Ortega-Sanchez, I. R., Broder, K. R., Talbot, H. K., Long, S. S., Havers, F. P., and Melgar, M. (2024). Use of respiratory syncytial virus vaccines in adults aged ≥60 years: Updated recommendations of the Advisory Committee on Immunization Practices—United States, 2024. *MMWR Morbidity and Mortality Weekly Report*, 73(32):696–702.

Bronars, S. G. (1987). The power of nonparametric tests of preference maximization. *Econometrica*, 55(3):693–698.

Bronchetti, E. T., Huffman, D. B., and Magenheim, E. (2015). Attention, intentions, and follow-through in preventive health behavior: Field experimental evidence on flu vaccination. *Journal of Economic Behavior & Organization*, 116:270–291.

Bruine de Bruin, W., Wallin, A., Parker, A. M., Strough, J., and Hanmer, J. (2017). Effects of anti- versus pro-vaccine narratives on responses by recipients varying in numeracy: A cross-sectional survey-based experiment. *Medical Decision Making*, 37(8):860–870.

Calvin, C. M., Deary, I. J., Fenton, C., Roberts, B. A., Der, G., Leckenby, N., and Batty, G. D. (2011). Intelligence in youth and all-cause-mortality: Systematic review with meta-analysis. *International Journal of Epidemiology*, 40(3):626–644.

Campos-Mercade, P., Meier, A. N., Schneider, F. H., and Wengström, E. (2021). Prosociality predicts health behaviors during the COVID-19 pandemic. *Journal of Public Economics*, 195:104367.

Cappelen, A. W., Kariv, S., Sørensen, E. Ø., and Tungodden, B. (2023). The development gap in economic rationality of future elites. *Games and Economic Behavior*, 142:866–878.

Carvalho, L., Olafsson, A., and Silverman, D. (2024). Misfortune and mistake: The financial conditions and decision-making ability of high-cost loan borrowers. *Journal of Political Economy*, 132(9):3173–3213.

Centers for Disease Control and Prevention (2018). Estimates of influenza vaccination coverage among adults—united states, 2017–18 flu season. FluVaxView, October 25.

Centers for Disease Control and Prevention (2020). Flu vaccination coverage, united states, 2019–20 influenza season. FluVaxView.

Centers for Disease Control and Prevention (2025). CDC immunization schedule adopts individual-based decision-making for COVID-19 and standalone vaccination for chickenpox in toddlers. News release, October 6.

Chen, M., Liu, T. X., Shan, Y., Wang, S., Zhong, S., and Zhou, Y. (2025). How general are measures of choice consistency? evidence from experimental and scanner data. Technical Report 2505.05275, arXiv. Version 2.

Choi, S., Fisman, R., Gale, D., and Kariv, S. (2007). Consistency and heterogeneity of individual behavior under uncertainty. *American Economic Review*, 97(5):1921–1938.

Choi, S., Kariv, S., Müller, W., and Silverman, D. (2014). Who is (more) rational? *American Economic Review*, 104(6):1518–1550.

Choi, Y. and Fox, A. M. (2022). Mistrust in public health institutions is a stronger predictor of vaccine hesitancy and uptake than trust in Trump. *Social Science & Medicine*, 314:115440.

Cinelli, M., Quattrociocchi, W., Galeazzi, A., Valensise, C. M., Brugnoli, E., Schmidt, A. L., Zola, P., Zollo, F., and Scala, A. (2020). The COVID-19 social media infodemic. *Scientific Reports*, 10:16598.

Cokely, E. T., Galesic, M., Schulz, E., Ghazal, S., and Garcia-Retamero, R. (2012). Measuring risk literacy: The Berlin Numeracy Test. *Judgment and Decision Making*, 7(1):25–47.

Cutler, D. M. and Lleras-Muney, A. (2010). Understanding differences in health behaviors by education. *Journal of Health Economics*, 29(1):1–28.

Dai, H., Saccardo, S., Han, M. A., Roh, L., Raja, N., Vangala, S., Modi, H., Pandya, S., Sloyan, M., and Croymans, D. M. (2021). Behavioural nudges increase COVID-19 vaccinations. *Nature*, 597(7876):404–409.

Dean, M. and Martin, D. (2016). Measuring rationality with the minimum cost of revealed preference violations. *The Review of Economics and Statistics*, 98(3):524–534.

Diesel, J., Sterrett, N., Dasgupta, S., Kriss, J. L., Barry, V., Vanden Esschert, K., Whiteman, A., Cadwell, B. L., Weller, D., Qualters, J. R., Harris, L., Bhatt, A., Williams, C., Fox, L. M., Meaney Delman, D., Black, C. L., and Barbour, K. E. (2021). COVID-19 vaccination coverage among adults—United States, December 14, 2020–May 22, 2021. *Morbidity and Mortality Weekly Report*, 70(25):922–927.

Elinder, M., Erixson, O., and Öhman, M. (2023). Cognitive ability, health policy, and the dynamics of COVID-19 vaccination. *Journal of Health Economics*, 91:102802.

Flannery, B., Chung, J. R., Belongia, E. A., et al. (2018). Interim estimates of 2017–18 seasonal influenza vaccine effectiveness—united states, february 2018. *MMWR Morbidity and Mortality Weekly Report*, 67(6):180–185.

Freeman, D., Loe, B. S., Yu, L.-M., Freeman, J., Chadwick, A., Vaccari, C., Shanyinde, M., Harris, V., Waite, F., Rosebrock, L., Petit, A., Vanderslott, S., Lewandowsky, S., Larkin, M., Innocenti, S., Pollard, A. J., McShane, H., and Lambe, S. (2021). Effects of different types of written vaccination information on COVID-19 vaccine hesitancy in the UK (OCEANS-III): A single-blind, parallel-group, randomised controlled trial. *The Lancet Public Health*, 6(6):e416–e427.

Gatz, M., Darling, J. E., Schneider, S., Liu, Y., Finkel, D., Orriens, B., Hernandez, R., Gutsche, T., Garcia, Y., and Kapteyn, A. (2026). An online population-representative longitudinal cognitive dataset from the Understanding America Study. *Scientific Data*, 13:698.

Gelbach, J. B. (2016). When do covariates matter? and which ones, and how much? *Journal of Labor Economics*, 34(2):509–543.

Geoffard, P.-Y. and Philipson, T. (1997). Disease eradication: Private versus public vaccination. *American Economic Review*, 87(1):222–230.

Gillen, B., Snowberg, E., and Yariv, L. (2019). Experimenting with measurement error: Techniques with applications to the Caltech Cohort Study. *Journal of Political Economy*, 127(4):1826–1863.

Grohskopf, L. A., Ferdinands, J. M., Blanton, L. H., Broder, K. R., and Loehr, J. (2024). Prevention and control of seasonal influenza with vaccines: Recommendations of the Advisory Committee on Immunization Practices – United States, 2024–25 influenza season. *MMWR Recommendations and Reports*, 73(5):1–25.

Halstead, I. N., McKay, R. T., and Lewis, G. J. (2022). COVID-19 and seasonal flu vaccination hesitancy: Links to personality and general intelligence in a large, UK cohort. *Vaccine*, 40(32):4488–4495.

Ho, L., Breza, E., Banerjee, A., Chandrasekhar, A. G., Stanford, F. C., Fior, R., Goldsmith-Pinkham, P., Holland, K., Hoppe, E., Jean, L.-M., Ogbu-Nwobodo, L., Olken, B. A., Torres, C., Vautrey, P.-L., Warner, E., Duflo, E., and Alsan, M. (2023). The impact of large-scale social media advertising campaigns on COVID-19 vaccination: Evidence from two randomized controlled trials. *AEA Papers and Proceedings*, 113:653–658.

Humlum, M. K., Morthorst, M. O., and Thingholm, P. R. (2024). Sibling spillovers and the choice to get vaccinated: Evidence from a regression discontinuity design. *Journal of Health Economics*, 94:102843.

IPUMS (2026). IPUMS NHIS: National health interview survey [dataset]. Minneapolis, MN: IPUMS. Data extract covering 2014–2022.

James, E. K., Bokemper, S. E., Gerber, A. S., Omer, S. B., and Huber, G. A. (2021). Persuasive messaging to increase COVID-19 vaccine uptake intentions. *Vaccine*, 39(49):7158–7165.

KFF (2026). KFF tracking poll on health information and trust: Trust in the CDC and views of federal childhood vaccine schedule changes. KFF. February 6. Survey conducted January 13–20, 2026.

Krupenkin, M. (2021). Does partisanship affect compliance with government recommendations? *Political Behavior*, 43(1):451–472.

LaCour, M. and Bell, Z. (2024). Attitudes towards COVID-19 vaccines may have 'spilled over' to other, unrelated vaccines along party lines in the United States. *Harvard Kennedy School Misinformation Review*, 5(3).

Larsen, B. J., Ryan, T. J., Greene, S., Hetherington, M. J., Maxwell, R., and Tadelis, S. (2023). Counter-stereotypical messaging and partisan cues: Moving the needle on vaccines in a polarized United States. *Science Advances*, 9(29):eadg9434.

Leiner, D. J. (2019). Too fast, too straight, too weird: Non-reactive indicators for meaningless data in internet surveys. *Survey Research Methods*, 13(3):229–248.

Li, J., Kariv, S., and Silverman, D. (2018). Economic rationality, health behaviors and health status. Conference presentation, 7th Annual Conference of the American Society of Health Economists.

Li, M., Taylor, E. G., Atkins, K. E., Chapman, G. B., and Galvani, A. P. (2016). Stimulating influenza vaccination via prosocial motives. *PLoS ONE*, 11(7):e0159780.

Loomba, S., de Figueiredo, A., Piatek, S. J., de Graaf, K., and Larson, H. J. (2021). Measuring the impact of COVID-19 vaccine misinformation on vaccination intent in the UK and USA. *Nature Human Behaviour*, 5:337–348.

Lunz Trujillo, K., Green, J., Safarpour, A., Lazer, D., Lin, J., and Motta, M. (2024). COVID-19 spillover effects onto general vaccine attitudes. *Public Opinion Quarterly*, 88(1):97–122.

Lusardi, A. and Mitchell, O. S. (2014). The economic importance of financial literacy: Theory and evidence. *Journal of Economic Literature*, 52(1):5–44.

MacNeil, J. R., Su, J. R., Broder, K. R., Guh, A. Y., Gargano, J. W., Wallace, M., Hadler, S. C., Scobie, H. M., Blain, A. E., Moulia, D., Daley, M. F., McNally, V. V., Romero, J. R., Talbot, H. K., Lee, G. M., Bell, B. P., and Oliver, S. E. (2021). Updated recommendations from the Advisory Committee on Immunization Practices for use of the Janssen (Johnson & Johnson) COVID-19 vaccine after reports of thrombosis with thrombocytopenia syndrome among vaccine recipients—United States, April 2021. *Morbidity and Mortality Weekly Report*, 70(17):651–656.

Melchinger, H., Belgaumi, S. M., Ahmed, N., Omer, S. B., and Malik, A. A. (2025). Change in influenza vaccine uptake among adults in the United States from May 2020 to October 2024. *PLOS Global Public Health*, 5(7):e0004756.

Melgar, M., Britton, A., Roper, L. E., Talbot, H. K., Long, S. S., Kotton, C. N., and Havers, F. P. (2023). Use of respiratory syncytial virus vaccines in older adults: Recommendations of the Advisory Committee on Immunization Practices—United States, 2023. *MMWR Morbidity and Mortality Weekly Report*, 72(29):793–801.

Milkman, K. L., Beshears, J., Choi, J. J., Laibson, D., and Madrian, B. C. (2011). Using implementation intentions prompts to enhance influenza vaccination rates. *Proceedings of the National Academy of Sciences*, 108(26):10415–10420.

Milkman, K. L., Ellis, S. F., Gromet, D. M., Jung, Y., Luscher, A. S., Mobarak, R. S., Paxson, M. K., Silvera Zumaran, R. A., Kuan, R., Berman, R., Lewis, Jr., N. A., List, J. A., Patel, M. S., Van den Bulte, C., Volpp, K. G., Beauvais, M. V., Bellows, J. K., Marandola, C. A., and Duckworth, A. L. (2024). Megastudy shows that reminders boost vaccination but adding free rides does not. *Nature*, 631:179–188.

Moehring, A., Collis, A., Garimella, K., Rahimian, M. A., Aral, S., and Eckles, D. (2023). Providing normative information increases intentions to accept a COVID-19 vaccine. *Nature Communications*, 14(1):126.

Nishimura, H., Ok, E. A., and Quah, J. K.-H. (2017). A comprehensive approach to revealed preference theory. *American Economic Review*, 107(4):1239–1263.

Panagiotakopoulos, L., Moulia, D. L., Godfrey, M., Link-Gelles, R., Roper, L., Havers, F. P., Taylor, C. A., Stokley, S., Talbot, H. K., Schechter, R., Brooks, O., Daley, M. F., Fleming-Dutra, K. E., and Wallace, M. (2024). Use of COVID-19 vaccines for persons aged ≥6 months: Recommendations of the Advisory Committee on Immunization Practices—United States, 2024–2025. *MMWR Morbidity and Mortality Weekly Report*, 73(37):819–824.

Pennycook, G., McPhetres, J., Bago, B., and Rand, D. G. (2022). Beliefs about COVID-19 in Canada, the United Kingdom, and the United States: A novel test of political polarization and motivated reasoning. *Personality and Social Psychology Bulletin*, 48(5):750–765.

Peters, E., Västfjäll, D., Slovic, P., Mertz, C. K., Mazzocco, K., and Dickert, S. (2006). Numeracy and decision making. *Psychological Science*, 17(5):407–413.

Philipson, T. (1995). The welfare loss of disease and the theory of taxation. *Journal of Health Economics*, 14(3):387–395.

Philipson, T. (1996). Private vaccination and public health: An empirical examination for U.S. measles. *Journal of Human Resources*, 31(3):611–630.

Philipson, T. (2000). Economic epidemiology and infectious diseases. In Culyer, A. J. and Newhouse, J. P., editors, *Handbook of Health Economics*, volume 1B, pages 1761–1799. Elsevier.

Pink, S. L., Chu, J., Druckman, J. N., Rand, D. G., and Willer, R. (2021). Elite party cues increase vaccination intentions among Republicans. *Proceedings of the National Academy of Sciences*, 118(32):e2106559118.

Polisson, M., Quah, J. K.-H., and Renou, L. (2020). Revealed preferences over risk and uncertainty. *American Economic Review*, 110(6):1782–1820.

Read, B., Wolters, L., and Berinsky, A. J. (2022). Racing the clock: Using response time as a proxy for attentiveness on self-administered surveys. *Political Analysis*, 30(4):550–569.

Reyna, V. F., Nelson, W. L., Han, P. K., and Dieckmann, N. F. (2009). How numeracy influences risk comprehension and medical decision making. *Psychological Bulletin*, 135(6):943–973.

Roozenbeek, J., Schneider, C. R., Dryhurst, S., Kerr, J., Freeman, A. L. J., Recchia, G., van der Bles, A. M., and van der Linden, S. (2020). Susceptibility to misinformation about COVID-19 around the world. *Royal Society Open Science*, 7(10):201199.

Sacks, D. W. and Sydnor, J. R. (2026). Preferences, beliefs, and demand for the flu vaccine. Working paper, University of Wisconsin–Madison. Revised September 23, 2026; originally circulated as NBER Working Paper No. 34230 (2025).

Sacks, D. W. and Walker, B. (2026). Can scientific information address vaccine hesitancy? a randomized controlled trial. Working paper, University of Wisconsin–Madison and Tulane University. July 13, 2026 version.

Sandlin, E. W. (2025). The politicization of influenza: Partisan changes in flu vaccination before and after COVID-19. *Journal of Public Health*, 47(2):317–325.

Schmid, P., Rauber, D., Betsch, C., Lidolt, G., and Denker, M.-L. (2017). Barriers of influenza vaccination intention and behavior: A systematic review of influenza vaccine hesitancy, 2005–2016. *PLoS ONE*, 12(1):e0170550.

Shim, E., Chapman, G. B., Townsend, J. P., and Galvani, A. P. (2012). The influence of altruism on influenza vaccination decisions. *Journal of the Royal Society Interface*, 9(74):2234–2243.

Solon, G., Haider, S. J., and Wooldridge, J. M. (2015). What are we weighting for? *Journal of Human Resources*, 50(2):301–316.

Varian, H. R. (1982). The nonparametric approach to demand analysis. *Econometrica*, 50(4):945–973.

Weller, J. A., Dieckmann, N. F., Tusler, M., Mertz, C. K., Burns, W. J., and Peters, E. (2013). Development and testing of an abbreviated numeracy scale: A Rasch analysis approach. *Journal of Behavioral Decision Making*, 26(2):198–212.

A Data and Samples

The analysis links four main Understanding America Study (UAS) sources by the respondent identifier `uasid`: the UAS Comprehensive Panel Dataset, which harmonizes UAS–HRS Waves 1–6; the main, by-round, and screen-timing files of the risky-budget experiments in UAS 5 and UAS 127; UAS 71, which records party identification; and Waves 21–30 of the UAS COVID-19 National Sample Longitudinal File. The UAS data are available to registered researchers at `https://uasdata.usc.edu`. The National Health Interview Survey (NHIS) benchmarks in Figure A1 use an IPUMS NHIS extract for 2014–2022 (IPUMS, 2026) and the 2023 and 2024 Sample Adult public-use files for influenza vaccination. The replication package lists the raw UAS fields used in the analysis. Derived variables, recodes, standardized measures, and period and phase indicators are documented in the replication code.

Every COVID-19 respondent is also in the influenza sample. The 575 respondents in the influenza sample only are younger, 59.6 against 62.9 years in 2021, and score lower on quantitative ability, −0.123 against 0.030 standard deviations, than the 2,364 respondents in both samples (Table A1); restricting the influenza analysis to the respondents in both samples gives a flu divergence of 2.00 (Figure 8, Panel A). Figure A1 compares UAS vaccination rates with the NHIS (Section 4.1). Additional subgroup and robustness characteristics are summarized in Table A2.

Table A1: Composition and Overlap of the Influenza and COVID-19 Analysis Samples

	(1) Influenza sample	(2) COVID-19 sample	(3) Influenza only	(4) Standardized difference
e^*	0.876	0.878	0.865	0.09
e^{**}	0.750	0.751	0.744	0.03
Quantitative ability (index)	0.00	0.03	−0.12	0.15
Age in 2021	62.22	62.87	59.56	0.24
Female	0.526	0.529	0.515	0.03
Asian	0.027	0.027	0.026	0.01
Hispanic	0.065	0.059	0.092	−0.13
Non-Hispanic Black	0.082	0.083	0.078	0.02
Non-Hispanic White	0.765	0.775	0.723	0.12
Other race	0.060	0.055	0.080	−0.10
Below high school	0.058	0.053	0.077	−0.10
Associate degree	0.138	0.146	0.104	0.13
Bachelor's degree	0.216	0.222	0.195	0.07
Graduate degree	0.165	0.169	0.150	0.05
High school	0.180	0.177	0.195	−0.05
Some college, no degree	0.243	0.234	0.280	−0.11

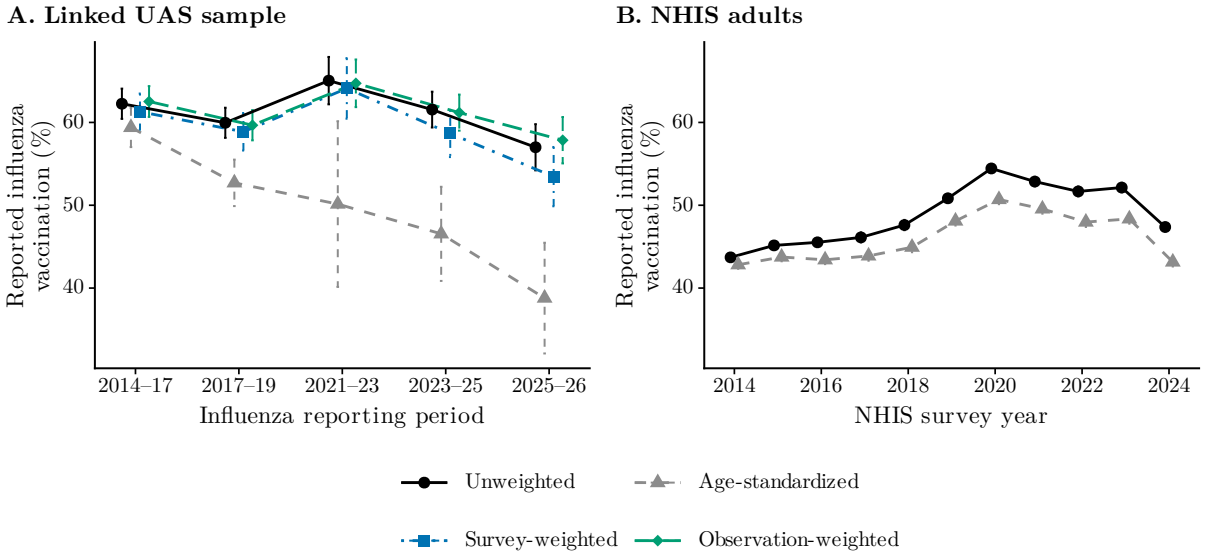
Notes: Data are from the Understanding America Study (UAS). Entries are respondent-level means. The influenza sample has 2,939 respondents and the COVID-19 sample 2,364. Every COVID-19 respondent is also in the influenza sample, so column 3 contains the remaining 575. Column 4 is the difference between columns 2 and 3 divided by the square root of the average of the two groups' variances. Means of binary variables are proportions.

Table A2: Characteristics Used in Subgroup and Robustness Analyses

	(1) Observations	(2) Mean	(3) SD	(4) Min	(5) Max
Additional respondent characteristics					
Democrat	2,208	0.357	0.479		
Republican	2,208	0.388	0.487		
Independent or other	2,208	0.255	0.436		
Born in 1955 or earlier	2,939	0.432	0.496		
Chronic conditions	2,936	1.24	1.24	0.00	7.00
Two or more conditions	2,936	0.362	0.481		
Household income ($1,000s)	2,903	115.40	149.56	0.00	2502.03
Health insurance	2,890	0.680	0.467		
Fair or poor health	2,939	0.185	0.388		

Notes: Data are from the Understanding America Study (UAS). Entries use the 2,939 respondents in the influenza sample with a nonmissing value of each characteristic. Party identification is from UAS 71 (November–December 2016). Respondents born in 1955 or earlier were 65 or older by the end of 2020. Chronic conditions count diabetes, high blood pressure, cancer, heart disease, lung disease, stroke, and arthritis. Chronic conditions, household income, health insurance, and self-rated health take each respondent's first nonmissing UAS–HRS value from Waves 1–2. Means of binary variables are proportions.

Figure A1: Influenza Vaccination Rates in the UAS and NHIS



Notes: Panel A uses the linked Understanding America Study (UAS) sample and plots the percentage of respondent-period observations with a reported flu shot: unweighted, age-standardized, weighted by UAS survey weights, and weighted by inverse-probability-of-observation weights. Panel B uses National Health Interview Survey (NHIS) adults aged 18 and older (IPUMS NHIS for 2014–2022 and the Sample Adult public-use files for 2023–2024) and plots the percentage vaccinated in the previous 12 months, unweighted and age-standardized. Age standardization uses the 2020 Census adult population in seven age groups (18–34, 35–44, 45–54, 55–64, 65–74, 75–84, and 85 or older). Each series uses the observations with nonmissing values for its weights or age. The UAS question covers about two seasons and the NHIS question one, and the samples differ, so the levels are not directly comparable. Bars are 95 percent confidence intervals, clustered by respondent for the UAS and treating NHIS observations as independent.

B Measures

Both indices measure how far a respondent's choices are from being rationalized by a well-behaved utility function, using Afriat's critical cost efficiency index (CCEI). The first, e^*, applies the CCEI to the observed choices; the second, e^{**}, applies it after adding mirror-image choices that impose symmetry across states and monotonicity with respect to first-order stochastic dominance (Choi et al., 2014).

B.1 Experimental Design

The experiment contained 25 scored allocation decisions under risk, preceded by two unscored practice rounds (Figure B1). In each round, the screen showed a budget line in a plane whose axes were dollar payoffs in two equally likely states, Heads and Tails. Respondents moved a cursor along the line, starting from a randomly selected feasible point, and submitted an allocation by clicking or pressing Enter; choices had to lie on the line. The Heads coordinate was paid if Heads occurred and the Tails coordinate if Tails occurred, so moving along the line traded payoffs across states at a relative price set by its slope. Budget lines were drawn independently across rounds and respondents; in every scored round, at least one intercept was \$50 or more and neither exceeded \$100. These limits apply to the scored rounds, not to the practice example in the figure. The data retain the continuous coordinates, although payments were rounded to cents.

After the 25 choices, the program selected one round at random and drew Heads or Tails with equal probability, and the respondent received the allocation to the realized state. No outcome feedback was given between rounds, so every choice could affect earnings and no realized payoff could affect later choices.

B.2 Formal Definitions

Appendix C.4 gives the revealed-preference relations and formal definitions of both indices.

B.3 Computation and Incomplete Records

A round enters the computation only if both intercepts and both coordinates are observed; missing rounds are not imputed. Of the 3,134 retained respondents, 3,066 have all 25 usable rounds and 68 have between 2 and 24; both indices require at least two rounds, which every retained respondent has. At each candidate e, the code builds the direct revealed-preference relation, takes its transitive closure, and applies the test in Equation (C12). A 36-step bisection on $[0, 1]$ locates each supremum to within 2×10^{-11}, using the T_i observations in D_i

for e_i^* and the $2T_i$ observations in D_i^M for e_i^{**}. When reporting shares at the upper bound, Table B1 classifies values within 10^{-8} of one as one; this rule does not enter the regressions.

Fewer usable rounds weaken the revealed-preference restrictions and cannot lower the index for a given choice record, so incomplete records can only raise it. Section 6.2 reports estimates for the 2,874 influenza and 2,313 COVID-19 respondents with all 25 rounds (Figure 8).

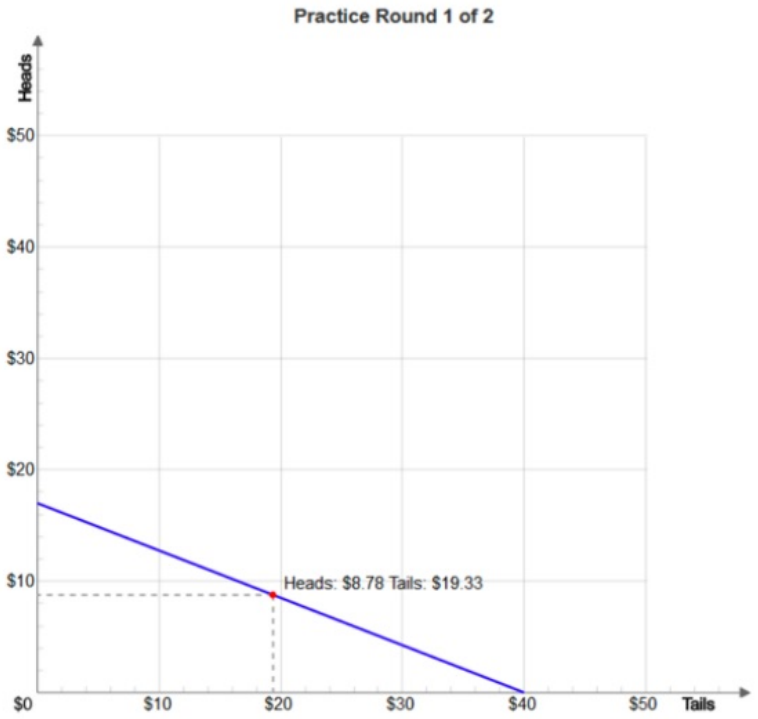
B.4 Random-Choice Benchmark

Following Bronars (1987), I compare the observed indices with those produced by random choice on the same budgets. For each module, I simulate 10,000 respondents. Each receives the 25-budget sequence of a randomly drawn respondent from that module and chooses a point uniformly along each budget line, and both indices are computed exactly as for the observed choices. The dashed curves in Figure 2, Panel A, show the simulated distributions. Depending on the module and index, 57–62 percent of observed values exceed the module-specific 95th percentile of the simulated distribution. Because the benchmark depends on the budget sequences, it measures the power of these budgets to detect violations rather than a fixed property of the indices (Choi et al., 2007, 2014). A value that does not exceed the benchmark does not by itself establish random choice or inattention.

B.5 Supplementary Statistics

Tables B2, B3, and B4 report the measure correlations, respondent-level correlates, and reliability statistics.

Figure B1: Risky-Budget Choice Interface



Notes: Screenshot of the first of two practice rounds in the UAS risky-budget task. Each point on the line pays the Heads amount or the Tails amount with probability one-half each, and respondents choose by moving the cursor along the line. Practice choices do not enter either consistency index. After the 25 scored rounds, the program draws one round for payment and then draws Heads or Tails with equal probability.

B.6 Quantitative-Ability Assessments

The UAS administers all four assessments repeatedly; the index uses each respondent's first score, from UAS–HRS Wave 1 or Wave 2 (Section 2.3). Three of the four have test-retest correlations of 0.74–0.77 between the two waves; serial sevens, whose mean is 4.57 out of 5, has 0.31 (Table B4). Table F2 reports the vaccination estimates when each component replaces the composite (Section 6.1).

Numeracy. Eight questions on probabilities and numerical information (Weller et al., 2013). Respondents compute the expected number of even outcomes in 1,000 rolls of a fair die; state the odds and the percentage chance that one ticket wins a 1,000-ticket lottery; convert a 10 percent disease rate into an expected case count and 20 out of 100 into a percentage; infer the probability of a tumor after a positive mammogram from a two-by-two table; and answer two cognitive-reflection items, the bat-and-ball and doubling-time problems. The UAS scores the responses with a two-parameter item-response model and reports a T-score with a mean of 50 and a standard deviation of 10 in a census-weighted sample of U.S. adults.

Number series. The Woodcock-Johnson Number Series test measures numerical pattern recognition: respondents supply the missing number in a sequence. The first administration uses 15 Form A items; the second uses 15 Form B items and four Form A anchor items, and item-response linking places both forms on a common scale. Public documentation gives examples such as 1, 2, __, 4 and 2, 4, 6, __; the scored sequences are licensed and not released. The UAS reports a T-score scaled as for numeracy.

Serial sevens. Respondents subtract seven from 100 and then from each answer, five subtractions in all; the score is the number of correct subtractions (0–5). The task draws on mental arithmetic, sustained attention, and working memory (Gatz et al., 2026). Most respondents reach the maximum; the mean is 4.57 (Table 1).

Financial literacy. Fourteen questions, scored as the number of correct answers (0–14) (Lusardi and Mitchell, 2014; Angrisani and Kapteyn, 2024). Five cover compound interest, inflation, purchasing power, and the timing of an inheritance; two cover the role of the stock market and mutual-fund diversification; and seven cover bond prices, relative risk and return, long-run asset performance, diversification, and housing-price risk.

Table B1: Additional Descriptive Statistics for Choice Consistency

	(1) Observations	(2) Mean	(3) SD	(4) Min	(5) Max
Panel A. Threshold shares and paired statistics					
$\Pr(e^* = 1)$	3,134	0.210			
$\Pr(e^* > 0.90)$	3,134	0.571			
$\Pr(e^* > 0.95)$	3,134	0.456			
$\Pr(e^{**} = 1)$	3,134	0.024			
$\Pr(e^{**} > 0.90)$	3,134	0.361			
$\Pr(e^{**} > 0.95)$	3,134	0.248			
e^* minus e^{**}	3,134	0.126	0.154	0.000	0.898
Pearson correlation: e^* and e^{**}	3,134	0.75			
Panel B. Statistics by experimental module					
UAS 5: e^*	1,461	0.879	0.146		
UAS 5: e^{**}	1,461	0.763	0.226		
UAS 5: e^* minus e^{**}	1,461	0.116	0.145		
UAS 127: e^*	1,673	0.873	0.150		
UAS 127: e^{**}	1,673	0.738	0.235		
UAS 127: e^* minus e^{**}	1,673	0.135	0.162		

	(1) Observations	(2) Mean
Panel C. Response quality: influenza respondents with timing data		
Fast responders	2,931	0.062
Slow responders	2,931	0.032
Failed interface task: fast	181	0.470
Failed interface task: others	2,731	0.152
Mean e**: fast	183	0.605
Mean e**: typical	2,653	0.758

Notes: Data are from the UAS choice experiments in Modules 5 and 127. Panels A and B use the 3,134 retained experiment respondents. Values within 10^{-8} of one count as equal to one, and the 0.90 and 0.95 thresholds are strict inequalities. The difference $e^* - e^{**}$ is the additional critical budget contraction needed when symmetry across states and monotonicity with respect to first-order stochastic dominance are imposed. Panel C uses respondents in the influenza sample with screen-timing data. The speed index averages, across screens, the ratio of the median time on that screen (within module and device type) to the respondent's time, capped at three (Leiner, 2019). Fast respondents have an index of at least 2, slow respondents an index below 0.5, and typical respondents an index of at least 0.5 but below 2. Failure rates use respondents with an observed result on the final paid task, which counts as passed when the selected y-coordinate is strictly between 31.49 and 32.51.

Table B2: Correlations among the Measures

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
(1) e^*									
(2) e^{**}	0.75								
(3) Quantitative ability	0.27	0.27							
(4) Numeracy	0.27	0.27	0.82						
(5) Number series	0.23	0.24	0.80	0.62					
(6) Serial sevens	0.10	0.09	0.59	0.24	0.27				
(7) Financial literacy	0.21	0.20	0.77	0.57	0.48	0.22			
(8) Years of education	0.15	0.17	0.46	0.42	0.35	0.18	0.43		
(9) Picture vocabulary	0.12	0.12	0.49	0.39	0.36	0.22	0.47	0.27	
(10) Verbal analogies	0.18	0.19	0.56	0.51	0.52	0.24	0.40	0.33	0.45

Notes: Data are from the Understanding America Study (UAS). Entries are Pearson correlations for the 2,893 respondents with every measure. Cognitive scores are each respondent's first pre-pandemic score. Years of education are approximated from the baseline education category.

Table B3: Demographic and Socioeconomic Correlates of Choice Consistency

	(1) e^*	(2) e^{**}
Birth year (per 10 years)	0.06***	0.08***
	(0.01)	(0.01)
Female	−0.17***	−0.16***
	(0.04)	(0.04)
Education: High school	0.04	0.12
	(0.09)	(0.09)
Education: Some college, no degree	0.17*	0.26***
	(0.09)	(0.09)
Education: Associate degree	0.19**	0.34***
	(0.10)	(0.09)
Education: Bachelor's degree	0.30***	0.42***
	(0.09)	(0.09)
Education: Graduate degree	0.44***	0.55***
	(0.09)	(0.09)
Household income (asinh, standardized)	0.06***	0.03*
	(0.02)	(0.02)
Respondents	2,958	2,958
Adjusted R^2	0.037	0.044

Notes: Entries are coefficients from respondent-level OLS regressions of each consistency index on the listed characteristics. Each index is standardized to mean zero and standard deviation one in this sample, and both columns use the same 2,958 respondents with nonmissing values of every regressor. Birth year is in units of 10 years. Education is the UAS–HRS Wave 1 category, or the Wave 2 category if Wave 1 is missing; below high school is the omitted category. Household income is the first nonmissing Wave 1–2 value, transformed by the inverse hyperbolic sine and standardized. Heteroskedasticity-robust standard errors are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table B4: Reliability of Choice Consistency and Quantitative Ability

	(1) UAS 5	(2) UAS 127	(3) Pooled
Panel A. Choice consistency: split-half statistics			
e^*			
Mean, odd rounds	0.94	0.93	0.94
Mean, even rounds	0.94	0.94	0.94
Share equal to one, odd rounds	0.481	0.452	0.465
Share equal to one, even rounds	0.522	0.500	0.510
Pearson correlation between halves	0.41	0.40	0.41
Spearman–Brown reliability	0.58	0.57	0.58
Correlation with quantitative ability, odd	0.18	0.23	0.21
Correlation with quantitative ability, even	0.20	0.19	0.19
e^{**}			
Mean, odd rounds	0.83	0.81	0.82
Mean, even rounds	0.83	0.81	0.82
Share equal to one, odd rounds	0.074	0.051	0.062
Share equal to one, even rounds	0.085	0.066	0.075
Pearson correlation between halves	0.58	0.60	0.59
Spearman–Brown reliability	0.73	0.75	0.75
Correlation with quantitative ability, odd	0.23	0.25	0.24
Correlation with quantitative ability, even	0.24	0.24	0.24
Respondents	1,422	1,663	3,085

Panel B. Quantitative ability and its components: Waves 1–2

	Correlation	Respondents
numeracy	0.77	2,568
number series	0.74	2,426
serial sevens	0.31	2,462
financial literacy	0.75	2,566
Quantitative-ability composite	0.82	2,344

Notes: Panel A recomputes each index on a respondent's odd-numbered and even-numbered scored rounds, for respondents with at least 12 usable rounds in each half. A half has 12 or 13 rounds and imposes fewer revealed-preference restrictions than the full set, so half-sample indices are higher on average and the correlations between halves may understate the reliability of the full indices. The Spearman–Brown reliability is $2r/(1 + r)$, where r is the Pearson correlation between halves. Correlations with quantitative ability use respondents in the construction sample. Because the experiment was administered once, Panel A does not measure stability over time. Panel B reports Pearson correlations between scores in UAS–HRS Wave 1 (2014–2017) and Wave 2 (2017–2019), for respondents with both scores. The composite requires all four components in both waves and applies the original construction-sample component means, standard deviations, and equal weights to each wave; its paired sample contains 2,344 respondents. The two number-series administrations use different item forms linked by item-response methods, so that row is an alternate-form correlation. Because the scores come from different waves, these correlations also reflect real change in skill.

B.7 COVID-19 Appointment and Reason Questions

Waves 25–29 asked these questions only of respondents who had not reported a dose or were unsure. Each of them was first asked whether they had tried to get a vaccine (`cr084`). Those who said yes were asked what they had tried (`cr085`), such as scheduling an appointment (`cr085s1`) or going to a walk-in site. Trying to schedule does not mean they got an appointment. People who tried to schedule then saw a list of booking problems (`cr086`), and people who tried a walk-in saw a separate list (`cr087`). People who said they had not tried, or were unsure, saw a list of reasons for not getting vaccinated (`cr088`). On each list they could choose more than one item.

I build three outcomes from these lists and do not use the walk-in list. The obstacle outcome is one if a person who tried to schedule said the appointment system was hard to understand (`cr086s4`), ran into technical problems (`cr086s5`), or missed a scheduled appointment (`cr086s6`). The how/where-to-book outcome is one if a person who did not report trying said she did not know how to schedule an appointment (`cr088s2`) or where to go to get vaccinated (`cr088s3`). The belief-reason outcome is one if she said the vaccine is not effective (`cr088s11`), not safe (`cr088s12`), or not necessary (`cr088s13`).

For Tables 5 and 6, I fix the cohort in Wave 21 and count each reason at most once per person, using only reports made before the first reported dose. I drop answers given after a person had already reported a dose. A zero means I see no report of that reason. This includes people who were vaccinated before the questions began, so a zero does not show that the person faced no barrier. Table 5 lets the three outcomes overlap, and Table 6 also shows each combination separately. None of the analyses condition on later beliefs, on willingness to be vaccinated, or on how many times a person was asked.

C Derivations

C.1 The Evaluation Probability and Its Derivatives

Condition on feasible vaccination and suppress the individual and date subscripts. Let $G \equiv G_{it}^F$, $c \equiv c_i > 0$, and $\sigma \equiv \sigma_t > 0$ denote the perceived net gain, decision precision, and decision difficulty, respectively, and let $z = cG/\sigma$. Equation (3) gives $P = \Phi(z)$ and $P_c = (G/\sigma)\phi(z)$, where the subscript denotes a partial derivative holding G and σ fixed. Differentiating the precision gradient with respect to σ gives

$$\frac{\partial P_c}{\partial \sigma} = -\frac{G}{\sigma^2}\,\phi(z) + \frac{G}{\sigma}\,\phi'(z)\left(-\frac{z}{\sigma}\right) = -\frac{G}{\sigma^2}\,\phi(z) + \frac{Gz^2}{\sigma^2}\,\phi(z) = \frac{G}{\sigma^2}\,\phi(z)\left(z^2 - 1\right). \tag{C1}$$

The second equality uses $\phi'(z) = -z\phi(z)$. For $G > 0$, the derivative is positive when $z > 1$ and negative when $0 < z < 1$. For a finite increase from σ_0 to $\sigma_1 > \sigma_0$, the precision gradient rises throughout the change if $cG/\sigma_1 > 1$.

For $G > 0$, the condition $z > 1$ is equivalent to $P = \Phi(z) > \Phi(1) \approx 0.84$. Greater decision difficulty therefore locally strengthens the precision gradient for an individual whose conditional vaccination probability exceeds this threshold. The condition concerns an individual probability given perceived gain, precision, difficulty, and feasibility; it is not a condition on the sample vaccination rate.

Changes in the perceived gain also alter the precision gradient. Holding c and σ fixed and using $\partial z/\partial G = c/\sigma$ gives

$$\frac{\partial P_c}{\partial G} = \frac{1}{\sigma}\,\phi(z) + \frac{G}{\sigma}\,\phi'(z)\,\frac{c}{\sigma} = \frac{1}{\sigma}\,\phi(z)\left(1 - z^2\right). \tag{C2}$$

By Equations (C1) and (C2), greater difficulty and a larger perceived gain need not increase the precision gradient. For an individual, the direction depends on the sign of G and on whether $|z|$ exceeds one; for the average gradient, it also depends on the distribution of perceived gains.

Responses to the environment. Section 3.2 also uses how the vaccination probability itself, rather than its precision gradient, responds to the environment. Holding c fixed,

$$\frac{\partial P}{\partial \sigma} = -\frac{z}{\sigma}\,\phi(z), \qquad \frac{\partial P}{\partial G} = \frac{c}{\sigma}\,\phi(z) = \frac{z}{G}\,\phi(z) \quad (G \neq 0). \tag{C3}$$

Among respondents with the same G and σ, both responses depend on $z\phi(z)$. This function is zero at $z = 0$, reaches its maximum $\phi(1) \approx 0.242$ at $z = 1$, and approaches zero as z grows. Respondents close to indifference and those with settled favorable assessments respond less

than respondents in between.

Consider two respondents with the same positive gain and $c_H > c_L$, so that $z_H > z_L$. For a change with $d\sigma \geq 0$ and $dG \leq 0$, not both zero,

$$d\left(P_H - P_L\right) = \left[z_L\phi(z_L) - z_H\phi(z_H)\right]\left(\frac{d\sigma}{\sigma} - \frac{dG}{G}\right). \tag{C4}$$

The second factor is positive, so the precision gradient widens when $z_H\phi(z_H) < z_L\phi(z_L)$. This condition holds whenever both z_L and z_H exceed one, that is, whenever both respondents vaccinate with probability above $\Phi(1) \approx 0.84$, because $z\phi(z)$ declines beyond $z = 1$. The lower-precision respondent's probability then falls more, because the higher-precision respondent lies farther along the declining portion of $z\phi(z)$. A more favorable change, with $d\sigma \leq 0$ and $dG \geq 0$, reverses the comparison, and when both z_L and z_H lie between zero and one, a less favorable change narrows the gradient. This between-respondent result is distinct from Equations (C1) and (C2), which describe how one respondent's own precision gradient changes. Equation (C4) is implication F2.

Risk ordering. Holding c_H, c_L, and σ fixed, both responses in Equation (C3) and the gap response in Equation (C4) approach zero as the positive gain G grows. A high-risk group with sufficiently large perceived gains can therefore respond less, and exhibit less widening, than a lower-risk group whose gains are smaller but still satisfy $c_L G/\sigma > 1$. This is a comparison with sufficiently settled favorable assessments, not a claim that widening decreases monotonically with perceived risk at every positive gain.

Two-season receipt. Section 3.2 maps the influenza outcome, any receipt over two seasonal decisions with independent disturbances, into $Q = 1-(1-P)^2$. Because $\partial Q/\partial c = 2(1-P)P_c$ has the sign of P_c, implication F1 carries over. For the comparison in Equation (C4), $dQ_H - dQ_L = 2(1-P_H)\,dP_H - 2(1-P_L)\,dP_L$. A less favorable change with $z_H\phi(z_H) < z_L\phi(z_L)$ gives $dP_L < dP_H \leq 0$, and $P_L < P_H$ gives $1 - P_L > 1 - P_H$, so $dQ_L < dQ_H$: the two-season gradient also widens through the decline at the bottom.

C.2 Learning and Access with a New Vaccine

Let B_i be the unknown expected net health benefit of vaccination, avoided illness loss less expected side-effect loss. Beliefs start from a normal prior with mean m_{i0} and precision $h_0 > 0$, common across individuals. By date t, individual i has received n_t conditionally independent signals $s_{i\ell} = B_i + \nu_{i\ell}$, with $\nu_{i\ell} \mid B_i \sim N(0, \eta_i^{-1})$, where $\eta_i > 0$ is the precision with which she extracts information from common evidence. Let b be a possible value of B_i and $\pi(b \mid s_{i1}, \ldots, s_{in_t})$ its posterior density. Bayes' rule multiplies the prior and signal

densities. Omitting factors that do not depend on b,

$$\pi(b \mid s_{i1}, \ldots, s_{in_t}) \propto \exp\left[-\frac{1}{2}\left\{h_0\,(b - m_{i0})^2 + \eta_i \sum_{\ell=1}^{n_t} (s_{i\ell} - b)^2\right\}\right]. \tag{C5}$$

Expanding the expression in braces and collecting terms in b gives

$$(h_0 + n_t\eta_i)\, b^2 - 2\left(h_0 m_{i0} + \eta_i \sum_{\ell=1}^{n_t} s_{i\ell}\right) b + C_0, \tag{C6}$$

where C_0 collects terms independent of b. Define $D_{it} = h_0 + n_t\eta_i$ and $M_{it} = h_0 m_{i0} + n_t\eta_i\bar{s}_{it}$, where $\bar{s}_{it} = n_t^{-1}\sum_{\ell=1}^{n_t} s_{i\ell}$ is her average processed signal. Completing the square rewrites Equation (C6) as $D_{it}(b - M_{it}/D_{it})^2 + C_1$, where C_1 also does not depend on b. Thus $\exp\{-\frac{1}{2}D_{it}(b - M_{it}/D_{it})^2\}$ is the kernel of a normal density with mean M_{it}/D_{it} and variance D_{it}^{-1}, and the posterior mean is

$$m_{it} = \frac{h_0\, m_{i0} + n_t\, \eta_i\, \bar{s}_{it}}{h_0 + n_t\, \eta_i}. \tag{C7}$$

Holding the prior and the observed signal average fixed, the quotient rule gives

$$\frac{\partial m_{it}}{\partial \eta_i} = \frac{n_t\bar{s}_{it}\,(h_0 + n_t\eta_i) - n_t\,(h_0 m_{i0} + n_t\eta_i\bar{s}_{it})}{(h_0 + n_t\eta_i)^2} = \frac{n_t\, h_0\,(\bar{s}_{it} - m_{i0})}{(h_0 + n_t\eta_i)^2}. \tag{C8}$$

More precise processing moves the assessment farther toward the evidence, upward when the evidence is more favorable than the prior and downward otherwise, so higher precision alone does not determine whether willingness rises or falls.

She accepts vaccination once $m_{it} \geq k_i$, the cost of obtaining a dose in the same units. Along a common favorable path of evidence, individuals share m_0, h_0, and k, and their signal averages equal s, with $m_0 < k < s$. Acceptance then requires $n_t \geq n_i^*$, where

$$n_i^* = \frac{h_0\,(k - m_0)}{\eta_i\,(s - k)}. \tag{C9}$$

Its comparative statics are

$$\begin{aligned} \frac{\partial n_i^*}{\partial \eta_i} &= -\frac{h_0(k - m_0)}{\eta_i^2(s - k)} < 0, & \frac{\partial n_i^*}{\partial m_0} &= -\frac{h_0}{\eta_i(s - k)} < 0, \\ \frac{\partial n_i^*}{\partial k} &= \frac{h_0(s - m_0)}{\eta_i(s - k)^2} > 0, & \frac{\partial n_i^*}{\partial s} &= -\frac{h_0(k - m_0)}{\eta_i(s - k)^2} < 0. \end{aligned} \tag{C10}$$

A more precise unfavorable prior also raises the threshold: $\partial n_i^*/\partial h_0 = (k - m_0)/[\eta_i(s - k)] > 0$.

Together with Equation (4), a nondecreasing signal count, a common access date, continuing availability, and no option value of waiting, these derivatives imply C1. For two individuals who both eventually pass their thresholds, the difference in receipt is zero before the earlier acceptance date, one between the two dates, and zero after the later one, so a receipt gap arising from information opens and then closes. With discrete signals, rounding the threshold up to an integer can equalize acceptance dates despite different thresholds, which is another reason the ranking in C1 is weak. The result conditions on the same favorable signal path; it is not an unconditional ranking of expected adoption times over all possible signal realizations.

The access-binding case, implication C2, follows directly from Equation (4). If $T_i^{\text{info}} < T_i^{\text{feas}}$ for all respondents above a given level of information precision, observed receipt cannot differ among them before access opens, because $T_i = T_i^{\text{feas}}$ regardless of when their information thresholds were crossed. Differences in receipt can appear once access becomes feasible. Differences in stated acceptance, by contrast, may be present earlier, because acceptance depends on T_i^{info} rather than on the access date.

C.3 Why a Narrowing Cumulative Gap Need Not Identify Learning

Let Q_t^g be cumulative receipt by date t in group g, and let h_t^g be the probability of first receipt at t among group members who were previously unvaccinated, so that $Q_t^g = Q_{t-1}^g + (1 - Q_{t-1}^g)\, h_t^g$. For groups H and L with the same subsequent hazard, $h_t^H = h_t^L = h_t$, subtraction gives

$$Q_t^H - Q_t^L = \left(Q_{t-1}^H - Q_{t-1}^L\right) + h_t \left[\left(1 - Q_{t-1}^H\right) - \left(1 - Q_{t-1}^L\right)\right] = (1 - h_t)\left(Q_{t-1}^H - Q_{t-1}^L\right). \quad \text{(C11)}$$

When $0 < h_t \leq 1$, equal subsequent hazards shrink an existing cumulative gap in absolute value without reversing its sign (implication C3). A narrowing cumulative gap therefore does not by itself indicate that differences in learning or acceptance have diminished. The argument concerns raw group differences with fixed membership and applies to the regression coefficients only by analogy.

C.4 Formal Definitions of the Consistency Indices

Respondent i contributes T_i usable scored choices, $D_i = \{(p_{ij}, x_{ij})\}_{j=1}^{T_i}$, where $x_{ij} = (x_{ij1}, x_{ij2})$ is the allocation chosen in round j. With axis intercepts b_{ij1} and b_{ij2}, prices are normalized as $p_{ij} = (1/b_{ij1}, 1/b_{ij2})$, so the budget is $\{x \in \mathbb{R}_+^2 : p_{ij} \cdot x \leq 1\}$ and a choice on the line satisfies $p_{ij} \cdot x_{ij} = 1$ up to numerical precision. Rescaling a round's prices does not

change its revealed-preference implications.

For a budget-retention factor $e \in [0,1]$, x_{ij} is e-directly revealed preferred to x_{ik} if $p_{ij} \cdot x_{ik} \leq e\, p_{ij} \cdot x_{ij}$, that is, if x_{ik} was affordable on round j's budget contracted by the factor e. Let $R_i(e)$ denote the transitive closure of this relation. The data satisfy e-GARP if, for all rounds j and k,

$$x_{ij}\, R_i(e)\, x_{ik} \implies p_{ik} \cdot x_{ij} \geq e\, p_{ik} \cdot x_{ik}, \tag{C12}$$

so a choice revealed preferred to another cannot lie strictly inside the other's contracted budget. Afriat's critical cost efficiency index (Afriat, 1973) is

$$e_i^* = \sup \{e \in [0,1] : D_i \text{ satisfies } e\text{-GARP}\}. \tag{C13}$$

Thus $1 - e_i^*$ is the infimum of the proportional budget contractions that remove every violation, and $e_i^* = 0.80$ corresponds to a 20 percent contraction. Because e_i^* is a supremum, $e_i^* = 1$ does not rule out violations exactly at the boundary. The index measures how coherent the choices are; it does not identify which round causes a violation or recover a utility function.

Heads and Tails are equally likely, and their labels carry no economic content. Following Choi et al. (2014), e^{**} adds the mirror image of each observation,

$$\tilde{p}_{ij} = (p_{ij2}, p_{ij1}), \qquad \tilde{x}_{ij} = (x_{ij2}, x_{ij1}), \qquad D_i^M = D_i \cup \{(\tilde{p}_{ij}, \tilde{x}_{ij})\}_{j=1}^{T_i}, \tag{C14}$$

and computes the CCEI of the augmented data:

$$e_i^{**} = \sup \left\{e \in [0,1] : D_i^M \text{ satisfies } e\text{-GARP}\right\}. \tag{C15}$$

In this two-state, equal-probability design, consistency of the augmented data requires, beyond ordinary GARP consistency, that choices be rationalizable by preferences that are symmetric across states and monotone with respect to first-order stochastic dominance (Choi et al., 2014; Nishimura et al., 2017; Polisson et al., 2020). Because $D_i \subseteq D_i^M$, any e that satisfies e-GARP for D_i^M also satisfies it for D_i, so $0 \leq e_i^{**} \leq e_i^* \leq 1$. The difference $e_i^* - e_i^{**}$ (Table B1) is the additional contraction needed to impose symmetry and dominance monotonicity jointly. It is not the share of dominated choices: a respondent whose choices fit stable but state-dependent preferences can have $e_i^{**} < e_i^*$ without choosing any dominated allocation.

Table C1 links these implications to the empirical comparisons.

Table C1: Model Implications and Empirical Comparisons

	(1) Prediction	(2) Empirical comparison	(3) Pattern that would contradict it
F1	An individual's precision gradient takes the sign of her perceived gain; its population average depends on the distribution of gains	Average influenza associations and their pre/post changes (Equations (5)–(6); Table 2); party-specific tertile changes (Figure 5; Appendix Table D4)	None without data on perceived gains; the party comparison is descriptive
F2	Under its stated conditions, a less favorable environment lowers vaccination more among less precise respondents who share a positive gain	Period-specific associations (Equation (7)); adjusted tertile paths (Figure 4); comparisons by age and chronic conditions (Figure 5; Appendix Table D4)	A widening driven by a rise at the top alone
C1	If information binds during the observed rollout, higher-ability respondents revise beliefs favorably faster and receive a dose weakly earlier	Wave-specific receipt associations (Equation (9); Figure 6); belief levels and changes (Figure 7; Appendix Table E2); beliefs added to the phase model (Table 4)	Earlier receipt without faster favorable belief revision during the period in which learning is assumed to generate the gap
Pre-rollout learning	If higher-ability respondents learn more from favorable evidence before access opens, their belief advantage is present at the start; their receipt gap opens only with access	Initial beliefs and later changes (Figure 7; Appendix Table E2); baseline-belief adjustment (Table 4)	Under a common prior and favorable evidence, no initial favorable belief advantage, or a receipt gap before access opens
C2	If access or arranging binds, the ability gradient in receipt opens with eligibility, even among respondents who accepted on the same date, and narrows as obtaining a dose becomes easier	Wave-specific receipt associations (Figure 6); reported barriers with and without Wave 21 beliefs (Table 5); reported reasons and later receipt (Table 6)	A receipt gradient that opens before eligibility does, or no difference in reported know-how or obstacles among respondents with the same Wave 21 beliefs
C3	Once both groups have accepted and access is open, equal positive subsequent hazards narrow a cumulative gap without any change in learning	Late-rollout wave associations (Figure 6)	No narrowing despite acceptance, open access, and equal positive hazards

Notes: F1 and F2 are the implications for a familiar vaccine and C1–C3 those for a new vaccine (Section 3 and Appendix C). The pre-rollout-learning row distinguishes learning before access opens from learning during the observed rollout. Column 2 lists the equations and exhibits that bear on each implication, and column 3 a pattern that would contradict it under its stated assumptions.

Table D1: Influenza Vaccination: Associations by Reporting Period

	(1) e^{**}	(2) Quantitative ability	(3) Both
e^{**}: 2014–2017	−1.52		−1.60
	(0.94)		(0.98)
2017–2019 minus 2014–2017	1.79**		1.72**
	(0.80)		(0.82)
2021–2023 minus 2014–2017	2.69*		2.38
	(1.48)		(1.50)
2023–2025 minus 2014–2017	2.61**		2.36**
	(1.17)		(1.19)
2025–2026 minus 2014–2017	4.48***		4.26***
	(1.45)		(1.47)
2021–2026 average minus 2017–2019	1.47		1.28
	(0.97)		(0.99)
Quantitative ability: 2014–2017		−0.42	−0.04
		(1.15)	(1.17)
2017–2019 minus 2014–2017		1.41	1.00
		(0.99)	(1.02)
2021–2023 minus 2014–2017		3.22*	2.67
		(1.95)	(1.97)
2023–2025 minus 2014–2017		2.75*	2.19
		(1.48)	(1.51)
2025–2026 minus 2014–2017		3.05	1.96
		(1.92)	(1.94)
2021–2026 average minus 2017–2019		1.60	1.28
		(1.27)	(1.29)
Observations	9,729	9,629	9,629
Respondents	3,038	2,939	2,939
Adjusted R^2	0.071	0.071	0.071

Notes: Entries are associations with reported influenza vaccination in 2014–2017 and changes from 2014–2017, in percentage points per standard deviation, from Equation (7). The last row in each block compares the equally weighted average of the three post-2020 changes with the 2017–2019 change. Columns 1 and 2 enter one measure at a time; column 3 enters both. Controls are those of Table 2, and every control coefficient except the one on months since the preceding interview differs by reporting period. e^{**} is scaled by its standard deviation across unique respondents in the estimation sample, and quantitative ability by its standard deviation in the construction sample. Standard errors clustered by respondent are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

D Influenza Supplementary Results

Table D1 reports period-specific continuous associations. Table D2 compares education coefficients on common samples. Table D3 reports joint-tertile paths. Figure D1 gives unadjusted rates. Table D4 reports the pooled post-2020 changes by group plotted in Figure 5. Measured skills account for little of the influenza education differences: adding both changes the 2014–2019 bachelor's-degree difference from 10.2 to 9.9 points and the graduate-degree difference from 16.3 to 16.0 (Table D2, Panel A, columns 1 and 4).

Table D2: Education Gaps in Vaccination with and without the Two Measures

	(1) No measure	(2) e^{**}	(3) Quantitative ability	(4) Both
Panel A. Influenza				
Bachelor's degree: 2014–2019	10.20***	10.49***	9.82**	9.90**
	(3.81)	(3.82)	(4.10)	(4.10)
2021–2026 minus 2014–2019	5.12	3.84	2.05	1.61
	(4.81)	(4.81)	(5.09)	(5.08)
Graduate's degree: 2014–2019	16.29***	16.67***	15.85***	16.01***
	(3.89)	(3.90)	(4.23)	(4.23)
2021–2026 minus 2014–2019	3.10	1.58	−0.36	−0.94
	(4.90)	(4.91)	(5.24)	(5.24)
Observations	9,629	9,629	9,629	9,629
Respondents	2,939	2,939	2,939	2,939
Adjusted R^2	0.072	0.072	0.072	0.072
Panel B. COVID-19				
Bachelor's degree: initial rollout	5.86***	5.73***	6.54***	6.46***
	(1.80)	(1.81)	(2.10)	(2.10)
Eligibility expansion minus initial	13.17***	13.95***	10.05***	10.48***
	(3.27)	(3.30)	(3.49)	(3.47)
Broad availability minus initial	20.42***	20.21***	13.28**	13.42**
	(5.13)	(5.18)	(5.30)	(5.31)
Late rollout minus initial	15.30***	15.34***	10.54**	10.70**
	(4.66)	(4.71)	(4.81)	(4.82)
Graduate's degree: initial rollout	8.34***	8.20***	9.11***	9.01***
	(1.95)	(1.97)	(2.27)	(2.27)
Eligibility expansion minus initial	17.49***	18.43***	14.00***	14.56***
	(3.40)	(3.44)	(3.68)	(3.67)
Broad availability minus initial	22.85***	22.61***	14.93***	15.09***
	(5.18)	(5.23)	(5.41)	(5.42)
Late rollout minus initial	16.08***	16.13***	10.74**	10.94**
	(4.72)	(4.76)	(4.92)	(4.93)
Observations	19,954	19,954	19,954	19,954
Respondents	2,364	2,364	2,364	2,364
Adjusted R^2	0.447	0.447	0.448	0.449

Notes: Entries are differences in vaccination between respondents in the listed education category and respondents without a high-school diploma, and changes in those differences, in percentage points. Panel A uses the influenza models of Table 2, Panel B (Equation (6)), and Panel B the COVID-19 models of Table 3 (Equation (8)). Column 1 omits both measures, columns 2 and 3 enter one measure, and column 4 enters both. Every model includes all six education categories; the high school, some college, and associate degree coefficients are not shown. Standard errors clustered by respondent are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table D3: Covariate-Adjusted Influenza Vaccination by Fixed Tertile

	(1) Bottom	(2) Middle	(3) Top	(4) Top minus bottom
Panel A. e^{**} tertiles				
2014–2017 level	63.79***	64.79***	59.02***	−4.77**
	(1.64)	(1.59)	(1.70)	(2.43)
2017–2019 minus 2014–2017	−4.27***	−5.49***	1.34	5.62***
	(1.36)	(1.38)	(1.41)	(2.01)
2021–2023 minus 2014–2017	−3.70	−0.69	3.70	7.41*
	(2.65)	(2.64)	(2.74)	(3.79)
2023–2025 minus 2014–2017	−4.35**	−1.26	1.64	5.99**
	(2.05)	(1.98)	(2.07)	(2.98)
2025–2026 minus 2014–2017	−9.36***	−5.25**	1.18	10.54***
	(2.56)	(2.54)	(2.53)	(3.65)
Observations	9,629	9,629	9,629	9,629
Respondents	2,939	2,939	2,939	2,939
Panel B. Quantitative ability tertiles				
2014–2017 level	62.04***	62.16***	63.34***	1.30
	(1.81)	(1.61)	(1.82)	(2.83)
2017–2019 minus 2014–2017	−3.61**	−2.55*	−2.20	1.41
	(1.60)	(1.31)	(1.41)	(2.31)
2021–2023 minus 2014–2017	−2.56	0.29	1.63	4.19
	(3.09)	(2.64)	(2.85)	(4.56)
2023–2025 minus 2014–2017	−2.54	−1.41	0.02	2.55
	(2.33)	(1.96)	(2.25)	(3.59)
2025–2026 minus 2014–2017	−4.20	−5.25**	−3.91	0.29
	(2.94)	(2.53)	(2.90)	(4.61)
Observations	9,629	9,629	9,629	9,629
Respondents	2,939	2,939	2,939	2,939

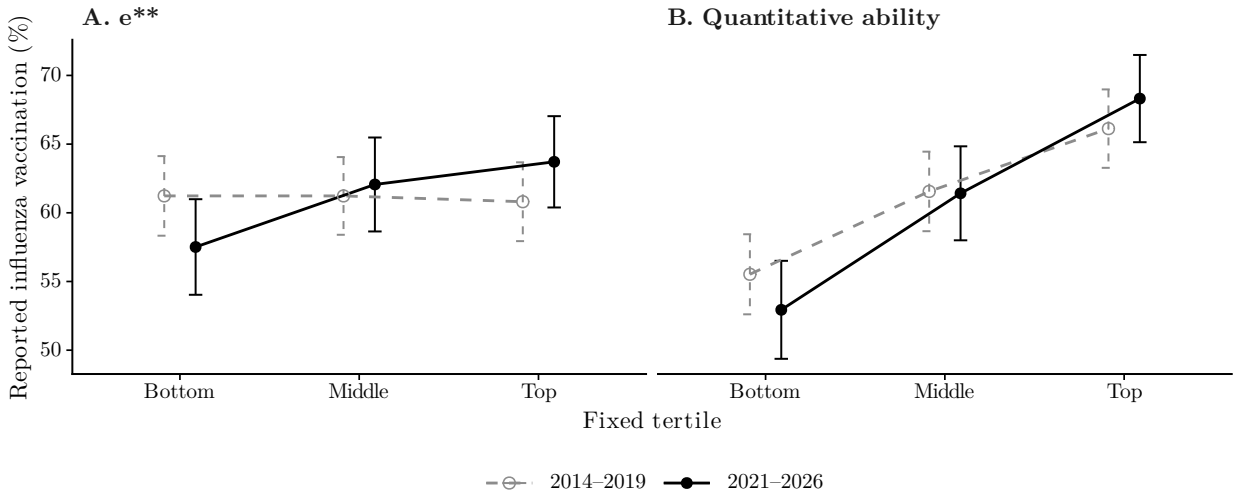
Notes: Panels A and B report the covariate-adjusted percentage reporting a flu shot in 2014–2017 and its change from 2014–2017, by fixed tertile of e^{**} and of quantitative ability. Both panels come from one linear probability model that interacts the e^{**} and quantitative-ability tertile indicators with reporting period. Sex, birth year, race and ethnicity, and education have period-specific coefficients, and months since the preceding interview has a single coefficient. Predictions are averaged over the covariates of the 2,939 respondents, with recall length at its mean. The tertile cutoffs are 0.6852 and 0.9129 for e^{**} and −0.3742 and 0.5257 for quantitative ability, and a value at a cutoff goes to the lower tertile. Standard errors clustered by respondent are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table D4: Change in Influenza Vaccination by Consistency Tertile within Groups

	Change in adjusted receipt (pp), 2014–2017 to 2021–2026			
	(1) Bottom e^{**} tertile	(2) Top e^{**} tertile	(3) Top minus bottom	(4) Respondents
All respondents	−5.62***	2.04	7.66***	2,939
	(1.87)	(1.92)	(2.72)	
By party				
Democrats	−0.76	3.19	3.95	788
	(3.38)	(3.61)	(5.06)	
Republicans	−15.27***	−4.74	10.53**	857
	(3.14)	(3.16)	(4.44)	
By birth cohort				
Born after 1955	−7.31***	1.59	8.89**	1,668
	(2.72)	(2.63)	(3.83)	
Born in 1955 or earlier	−3.77	1.84	5.62	1,271
	(2.48)	(2.71)	(3.73)	
By chronic conditions				
Fewer than two	−7.09***	4.34*	11.43***	1,874
	(2.51)	(2.35)	(3.49)	
Two or more	−3.97	−2.74	1.23	1,062
	(2.67)	(3.32)	(4.26)	

Notes: Entries are covariate-adjusted changes in the percentage reporting a flu shot from 2014–2017 to 2021–2026, in percentage points, by fixed e^{**} tertile. Each row comes from a separate model for that group, which interacts e^{**} and quantitative-ability tertile indicators with reporting period and includes sex, birth year, race and ethnicity, and education with period-specific coefficients and months since the preceding interview with a single coefficient. The three post-2020 reporting periods are pooled using their shares of observed respondent-period records; Appendix Table D3 reports each period separately for all respondents. Predictions are averaged over the group's respondents, holding their covariates fixed, and tertile cutoffs are the full-sample values. The party rows include Democrats and Republicans from UAS 71 (2016). Respondents born in 1955 or earlier were 65 or older by the end of 2020. Chronic conditions are the baseline count of diabetes, high blood pressure, cancer, heart disease, lung disease, stroke, and arthritis. Figure 5 plots the paths by period. Standard errors clustered by respondent are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Figure D1: Unadjusted Influenza Vaccination by Fixed Tertile

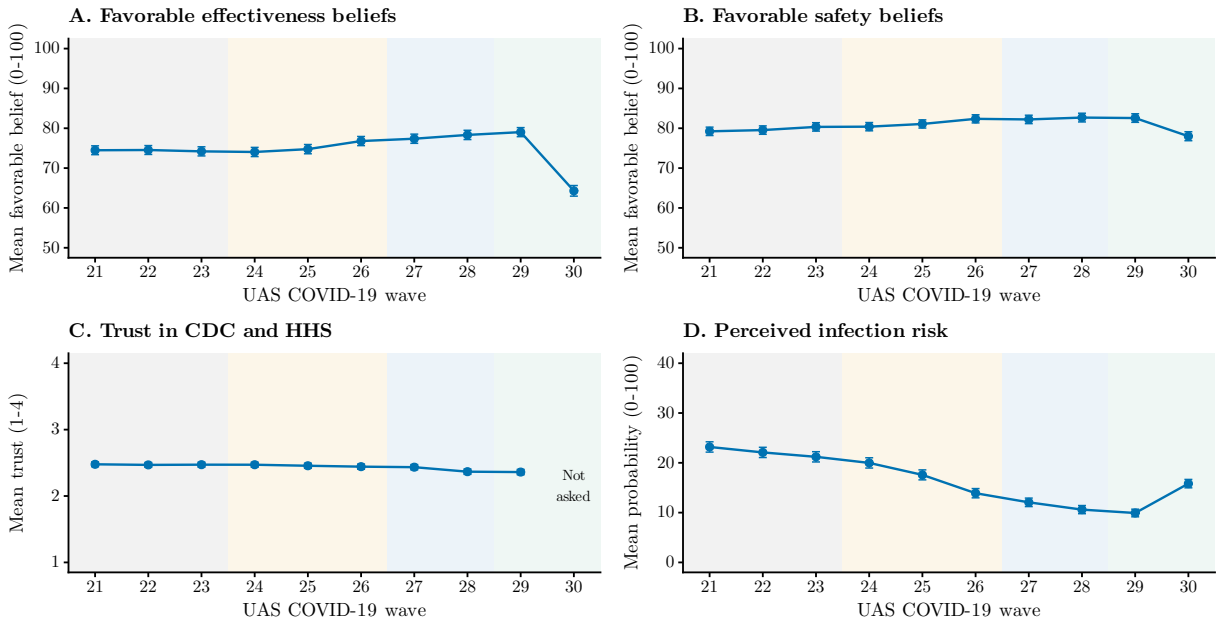


Notes: The figure plots the unadjusted percentage of observations with a reported flu shot by fixed tertile of e^{**} (Panel A) and of quantitative ability (Panel B), separately for 2014–2019 and 2021–2026. Tertile cutoffs are those in Appendix Table D3. The percentages do not adjust for covariates or for the other measure. Panel A uses 9,729 respondent-period observations from 3,038 respondents and Panel B 9,629 from 2,939. Bars are 95 percent confidence intervals clustered by respondent.

E COVID-19 Supplementary Results

Table E1 reports continuous wave-specific receipt associations. Figure E1 plots unadjusted wave means of vaccine beliefs, institutional trust, and perceived infection risk, and Table E2 reports their associations with e^{**} and quantitative ability by wave.

Figure E1: COVID-19 Vaccine Beliefs, Institutional Trust, and Perceived Infection Risk by Survey Wave



Notes: Each point is the unweighted mean among respondents with an observed item in that wave, in the COVID-19 receipt sample (19,954 respondent-wave observations from 2,364 respondents). Favorable effectiveness and safety beliefs equal 100 minus the reported probability that a vaccinated person catches COVID-19 or that vaccination causes serious side effects. Trust averages the CDC and HHS items (each 1–4) and requires both; it was not asked in Wave 30. Perceived infection risk is the stated probability (0–100) of infection in the next three months. These are descriptive means; the respondents observed can differ across waves, so changes need not be within-person changes. Shading marks the four rollout phases. Bars are pointwise 95 percent t confidence intervals for each mean.

Table E1: COVID-19 Receipt: Associations by Survey Wave

	(1) e^{**}	(2) Quantitative ability	(3) Both
e^{**}: Wave 21	0.42		0.45*
	(0.26)		(0.25)
Wave 22 minus Wave 21	−0.08		−0.11
	(0.50)		(0.51)
Wave 23 minus Wave 21	−0.48		−0.27
	(0.79)		(0.80)
Wave 24 minus Wave 21	−1.68*		−1.81*
	(0.97)		(0.97)
Wave 25 minus Wave 21	−1.03		−1.41
	(1.05)		(1.07)
Wave 26 minus Wave 21	−1.85*		−2.45**
	(1.08)		(1.09)
Wave 27 minus Wave 21	−0.02		−0.79
	(1.02)		(1.04)
Wave 28 minus Wave 21	0.26		−0.51
	(1.01)		(1.03)
Wave 29 minus Wave 21	0.45		−0.22
	(1.00)		(1.02)
Wave 30 minus Wave 21	−0.89		−1.26
	(0.88)		(0.89)
Quantitative ability: Wave 21		−0.06	−0.18
		(0.40)	(0.40)
Wave 22 minus Wave 21		0.15	0.18
		(0.64)	(0.66)
Wave 23 minus Wave 21		−1.39	−1.32
		(0.95)	(0.96)
Wave 24 minus Wave 21		0.39	0.82
		(1.13)	(1.14)
Wave 25 minus Wave 21		1.98	2.33*
		(1.27)	(1.29)
Wave 26 minus Wave 21		3.08**	3.70***
		(1.37)	(1.40)
Wave 27 minus Wave 21		4.65***	4.84***
		(1.30)	(1.33)
Wave 28 minus Wave 21		4.67***	4.80***
		(1.27)	(1.30)
Wave 29 minus Wave 21		4.11***	4.17***
		(1.25)	(1.28)
Wave 30 minus Wave 21		2.06*	2.37**
		(1.10)	(1.12)
Ability: Wave 30 minus 27, p-value		0.021	0.028
Observations	19,954	19,954	19,954
Respondents	2,364	2,364	2,364
Adjusted R^2	0.444	0.446	0.446

Notes: Entries are associations with cumulative self-reported receipt of at least one COVID-19 dose in Wave 21 and changes from Wave 21, in percentage points per standard deviation, from Equation (9). Columns 1 and 2 enter one measure at a time; column 3 enters both. All models include wave fixed effects and wave-specific coefficients on sex, birth year, race and ethnicity, education, baseline state, and interview date. e^{**} is scaled by its standard deviation across unique respondents in the estimation sample, and quantitative ability by its standard deviation in the construction sample. The p-value tests whether the Wave 30 and Wave 27 changes in the quantitative-ability association are equal. Standard errors clustered by respondent are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table E2: COVID-19 Vaccine Views: Associations by Survey Wave

	Panel A. Effectiveness		*Panel B. Safety*	
	(1) e^{**}	(2) Quantitative ability	(3) e^{**}	(4) Quantitative ability
Wave 21	0.16 (0.56)	4.75*** (0.74)	0.40 (0.51)	5.60*** (0.67)
Wave 22	-0.84 (0.55)	4.60*** (0.73)	-0.04 (0.52)	6.01*** (0.67)
Wave 23	-0.24 (0.58)	4.78*** (0.70)	-0.03 (0.52)	6.06*** (0.64)
Wave 24	-0.27 (0.58)	5.24*** (0.73)	0.48 (0.51)	6.16*** (0.68)
Wave 25	-0.28 (0.61)	4.37*** (0.77)	-0.04 (0.51)	6.15*** (0.66)
Wave 26	-0.57 (0.57)	3.53*** (0.79)	0.00 (0.50)	4.84*** (0.68)
Wave 27	-0.78 (0.55)	4.86*** (0.74)	0.17 (0.52)	6.47*** (0.69)
Wave 28	-0.06 (0.57)	4.94*** (0.74)	0.14 (0.54)	5.36*** (0.71)
Wave 29	0.09 (0.56)	5.00*** (0.71)	0.53 (0.55)	6.09*** (0.69)
Wave 30	-0.45 (0.66)	4.77*** (0.85)	-0.23 (0.54)	7.04*** (0.73)
Wave 27 minus Wave 21	-0.93 (0.61)	0.11 (0.78)	-0.23 (0.52)	0.87 (0.66)
Observations	20,167		20,157	
Respondents	2,377		2,377	
	Panel C. CDC/HHS trust		*Panel D. Infection risk*	
	(1) e^{**}	(2) Quantitative ability	(3) e^{**}	(4) Quantitative ability
Wave 21	0.040* (0.024)	0.063** (0.028)	-0.34 (0.57)	-1.96*** (0.73)
Wave 22	0.009 (0.024)	0.081*** (0.029)	-0.39 (0.56)	-2.11*** (0.71)
Wave 23	0.044* (0.024)	0.055* (0.029)	0.62 (0.53)	-2.84*** (0.73)
Wave 24	0.042* (0.025)	0.070** (0.029)	0.43 (0.54)	-2.97*** (0.75)
Wave 25	0.020 (0.025)	0.072** (0.029)	0.24 (0.54)	-2.82*** (0.72)
Wave 26	0.030 (0.025)	0.079** (0.031)	0.89* (0.48)	-3.58*** (0.69)
Wave 27	0.010 (0.026)	0.085*** (0.032)	0.18 (0.43)	-2.99*** (0.62)
Wave 28	0.032 (0.027)	0.078** (0.031)	-0.26 (0.44)	-3.81*** (0.64)
Wave 29	0.016 (0.027)	0.064** (0.031)	-0.13 (0.40)	-2.72*** (0.56)
Wave 30	–	–	0.13 (0.44)	-2.33*** (0.60)
Wave 27 minus Wave 21	-0.030 (0.021)	0.022 (0.025)	0.52 (0.59)	-1.04 (0.70)
Observations	18,032		20,182	
Respondents	2,305		2,378	

Notes: Entries are total associations of e^{**} and quantitative ability with each outcome in each wave from joint models; the contrast is the change from Wave 21 to Wave 27. Favorable effectiveness and safety beliefs are 100 minus the stated probability that a vaccinated person will contract COVID-19 or that vaccination will cause serious side effects. Perceived infection risk is the stated probability of infection in the next three months (0–100), without reversal. These three outcomes are in percentage points per score standard deviation. CDC/HHS trust requires both items, averages their 1–4 scores, and is standardized using its Wave 21 mean and standard deviation; coefficients are in Wave 21 trust standard deviations per score standard deviation. Trust was not asked in Wave 30. Quantitative ability retains its construction-sample standardization; e^{**} uses its standard deviation across unique respondents in each outcome's sample. Each model includes wave fixed effects and wave-specific coefficients on both measures, sex, birth year, race and ethnicity, education, baseline state, and interview date. The contrast compares associations rather than individual belief revisions. Observations are respondent-wave pairs. Standard errors clustered by respondent are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

F Robustness

Table F1 retains the e^* specifications. Table F2 replaces quantitative ability with its individual components. Table F3 reports observation and item-missingness associations. Table F4 reports the principal sensitivity checks displayed in Figure 8 and their same-sample baselines. For influenza observation weights, the numerator is the reporting-period observation rate and the denominator is a logit observation probability allowing coefficients on demographics, education, both consistency indices and quantitative ability to differ by period. The observation model includes potential records before the May 2022 restriction because missing interviews have no interview date. For COVID-19, the numerator conditions on wave and prior-wave observation and the denominator additionally allows baseline demographics, education and both measures to vary by wave. Predictions are bounded to [0.01,0.99]; numerator predictions are divided by denominator predictions. Observation and survey weights are not combined or normalized, and vaccination regressions treat estimated weights as fixed. Response-timing checks omit records missing timing data.

Table F1: Main Estimates Using e^*

	(1) e^* only	(2) e^* + quantitative ability
Panel A. Influenza		
e^*: average association	0.69	0.47
	(0.78)	(0.81)
e^*: 2014–2019	−0.11	−0.26
	(0.83)	(0.86)
2021–2026 minus 2014–2019	1.83*	1.62
	(0.98)	(1.00)
Quantitative ability: average association		1.09
		(0.95)
Quantitative ability: 2014–2019		0.34
		(1.04)
2021–2026 minus 2014–2019		1.85
		(1.21)
Observations	9,729	9,629
Respondents	3,038	2,939
Adjusted R^2	0.072	0.072
Panel B. COVID-19		
e^*: initial rollout	0.09	0.18
	(0.40)	(0.40)
Eligibility expansion minus initial	−0.77	−1.19
	(0.72)	(0.73)
Broad availability minus initial	0.52	−0.35
	(0.95)	(0.97)
Late rollout minus initial	0.54	−0.02
	(0.87)	(0.88)
Quantitative ability: initial rollout		−0.53
		(0.58)
Eligibility expansion minus initial		2.53***
		(0.86)
Broad availability minus initial		5.19***
		(1.21)
Late rollout minus initial		3.43***
		(1.09)
Observations	19,954	19,954
Respondents	2,364	2,364
Adjusted R^2	0.447	0.448
Panel C. e^* tertiles: top-minus-bottom gap change		
2017–2019 minus 2014–2017	6.36***	
	(1.96)	
2021–2023 minus 2014–2017	5.00	
	(3.78)	
2023–2025 minus 2014–2017	3.61	
	(2.95)	
2025–2026 minus 2014–2017	9.68***	
	(3.62)	
Observations	9,629	
Respondents	2,939	
Adjusted R^2	0.070	

Notes: Panels A and B repeat the models of Tables 2 and 3 with e^* in place of e^{**}, using the same controls and scaling. In Panel A, the average association comes from Equation (5) and the other rows from Equation (6); Panel B uses Equation (8). Panel C reports the change since 2014–2017 in the covariate-adjusted top-minus-bottom e^* tertile gap, in percentage points, from the model of Appendix Table D3 with e^* tertiles in place of e^{**} tertiles. Standard errors clustered by respondent are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table F2: Quantitative-Ability Components and Vaccination

	(1) Four-test index	(2) Numeracy	(3) Number series	(4) Serial sevens	(5) Financial literacy
Panel A. Influenza, conditional on e^{**}					
Ability: 2014–2019	0.44	0.68	−0.42	−0.43	1.64
	(1.04)	(1.00)	(0.95)	(0.86)	(1.02)
2021–2026 minus 2014–2019	1.78	2.17*	1.19	1.08	−0.18
	(1.21)	(1.15)	(1.05)	(0.98)	(1.20)
e^{**}: 2021–2026 minus 2014–2019	2.01**	2.00**	2.11**	2.22**	2.32**
	(0.95)	(0.95)	(0.94)	(0.94)	(0.95)
Observations	9,629	9,629	9,629	9,629	9,629
Respondents	2,939	2,939	2,939	2,939	2,939
Adjusted R^2	0.072	0.073	0.072	0.072	0.072
Panel B. COVID-19, conditional on e^{**}					
Ability: initial rollout	−0.55	−0.29	−1.15**	−0.25	0.46
	(0.58)	(0.58)	(0.52)	(0.44)	(0.53)
Eligibility expansion minus initial	2.65***	1.70**	1.31	1.12	2.52***
	(0.86)	(0.83)	(0.80)	(0.70)	(0.83)
Broad availability minus initial	5.22***	3.43***	2.83***	2.86***	3.52***
	(1.21)	(1.13)	(1.07)	(0.97)	(1.20)
Late rollout minus initial	3.57***	3.08***	2.08**	1.57*	2.14**
	(1.09)	(1.05)	(0.97)	(0.87)	(1.07)
e^{**}: expansion minus initial	−1.75**	−1.58**	−1.52**	−1.40*	−1.59**
	(0.74)	(0.74)	(0.74)	(0.73)	(0.73)
Observations	19,954	19,954	19,954	19,954	19,954
Respondents	2,364	2,364	2,364	2,364	2,364
Adjusted R^2	0.449	0.448	0.447	0.448	0.449

Notes: Each column enters one ability measure jointly with e^{**}: the four-test index in column 1 and one component in each of columns 2–5. Panel A follows Equation (6) and Panel B Equation (8), with the controls of Tables 2 and 3. Each component is standardized in the construction sample, and entries are in percentage points per standard deviation. The components are correlated, so these separate models do not isolate independent channels. Standard errors clustered by respondent are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table F3: Choice Consistency, Quantitative Ability, and Missing Influenza Reports

	(1)	(2)	(3)	(4)	(5)	(6)
	e^*		e^{**}		Quantitative ability	
	Absent from period	Item missing	Absent from period	Item missing	Absent from period	Item missing
Measure (per SD)	−0.75	0.01	−0.55	−0.01	−1.57***	0.04*
	(0.52)	(0.02)	(0.52)	(0.02)	(0.57)	(0.02)
Observations	15,655	10,989	15,655	10,989	14,695	10,889
Respondents	3,131	3,038	3,131	3,038	2,939	2,939
Mean of outcome (%)	29.8	0.1	29.8	0.1	25.9	0.1
Adjusted R^2	0.153	−0.001	0.153	−0.001	0.177	−0.001

Notes: Entries are associations of each measure with a missing influenza report, in percentage points per standard deviation, from unweighted linear probability models. In the absence columns, each respondent has one record for each of the five reporting periods, and the outcome is one when no interview is recorded in that period. In the item-missing columns, the sample is completed interviews, and the outcome is one when the flu-shot question lacks a valid yes or no. Records are built before the May 2022 interview-date restriction because records without an interview have no date. All models include reporting-period indicators, sex, birth year, race and ethnicity, and six education categories. Consistency indices are scaled by their standard deviation across unique respondents in the estimation sample, and quantitative ability by its standard deviation in the construction sample. Standard errors clustered by respondent are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table F4: Sensitivity of the Flu Divergence and the Availability Widening

	(1)	(2)	(3)	(4)	(5)	(6)
	Flu divergence (e^{**})		Availability widening			
	Estimate (s.e.)	Same-sample baseline	Estimate (s.e.)	Same-sample baseline	Flu resp.	COVID-19 resp.
Main specification	2.01 (0.95)	–	5.22 (1.21)	–	2,939	2,364
Measurement						
Each measure alone	2.25 (0.93)	–	5.10 (1.19)	–	3,038	2,364
All 25 rounds	1.74 (0.95)	–	5.03 (1.22)	–	2,874	2,313
Module trends	2.01 (0.95)	2.01	5.22 (1.21)	5.22	2,939	2,364
IV: split-half e^{**}	2.84 (1.69)	1.73	5.33 (1.24)	5.25	2,893	2,328
IV: retest ability	1.89 (1.01)	2.19	8.35 (1.84)	5.62	2,343	1,991
Response quality						
Passed interface check	1.81 (1.00)	–	5.35 (1.41)	–	2,419	1,955
Speed indicators	2.15 (0.96)	2.01	5.12 (1.21)	5.31	2,931	2,357
Exclude fast (index ≥ 2)	1.60 (0.99)	–	4.61 (1.26)	–	2,748	2,220
Exclude fast (index ≥ 1.5)	0.93 (1.07)	–	4.76 (1.35)	–	2,314	1,889
Log screen time	2.23 (0.97)	2.01	5.27 (1.21)	5.31	2,931	2,357
Other traits						
Party identification	2.15 (1.05)	2.28	5.53 (1.40)	5.30	2,208	1,809
Cheaper-state share	2.02 (0.95)	2.01	4.95 (1.22)	5.22	2,939	2,364
Income, health, insurance	2.02 (0.95)	2.05	5.22 (1.23)	5.25	2,887	2,339
Big Five	2.11 (0.95)	2.01	5.62 (1.22)	5.17	2,934	2,361
Weights and samples						
Survey weights	2.09 (1.17)	–	4.87 (1.66)	–	2,848	2,364
Observation weights	1.98 (0.96)	–	5.16 (1.22)	–	2,939	2,364
Balanced panel	2.20 (1.34)	–	5.90 (1.35)	–	784	1,552
Both samples (flu only)	2.00 (0.97)	–	–	–	2,364	–

Notes: Entries are the flu divergence, the change in the e^{**} association with reported influenza vaccination from 2014–2019 to 2021–2026 (Equation (6)), and the availability widening, the change in the quantitative-ability association with COVID-19 receipt from the initial rollout to broad adult availability (Equation (8)), in percentage points per standard deviation. Models enter both measures jointly except in the single-measure row. Columns 2 and 4 report the corresponding OLS estimates on the same observations for rows adding controls or instruments; sample restrictions and weighting rows have no separate baseline. Module controls retain the full main sample. The split-half IV row stacks the odd/even and even/odd directions and clusters by respondent; its OLS baseline uses the full consistency index. The ability-retest IV and its OLS baseline both use the Wave 1 composite; the instrument is its Wave 2 counterpart, with complete components in both waves. Respondent counts are unique people, not stacked records. Other rows are defined in Figure 8. Added controls have period- or phase-specific coefficients. Standard errors clustered by respondent are in parentheses. Entries retain the precision of the underlying replication outputs.